

## Eyes in Motion: A New Traffic Sensing Paradigm for Pedestrians and Cyclists

Vial, A.A.

**DOI**

[10.4233/uuid:e9cc7cc5-227c-464c-8d46-cb1b4073b7a2](https://doi.org/10.4233/uuid:e9cc7cc5-227c-464c-8d46-cb1b4073b7a2)

**Publication date**

2025

**Document Version**

Final published version

**Citation (APA)**

Vial, A. A. (2025). *Eyes in Motion: A New Traffic Sensing Paradigm for Pedestrians and Cyclists*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:e9cc7cc5-227c-464c-8d46-cb1b4073b7a2>

**Important note**

To cite this publication, please use the final published version (if applicable).  
Please check the document version above.

**Copyright**

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

**Takedown policy**

Please contact us and provide details if you believe this document breaches copyrights.  
We will remove access to the work immediately and investigate your claim.

The background is a light gray with large, dark gray letters 'A', 'M', 'P', 'S' at the top and 'A', 'M', 'P', 'S' at the bottom. A network of red lines connects various points, including circles and squares, across the page. Some lines are solid, while others are dashed. The overall aesthetic is modern and technical.

# **Eyes in Motion:** **A New Traffic Sensing Paradigm for** **Pedestrians and Cyclists**

**Alphonse Vial**

# **Eyes in Motion**

## **A New Traffic Sensing Paradigm for Pedestrians and Cyclists**

Alphonse VIAL

Delft University of Technology



The research leading to this dissertation has received funding from the European Research Council under the European Union Horizon 2020 Framework Programme for Research and Innovation. It is established by the Scientific Council of the ERC Grant Agreement no. 669792 (Allegro).

Cover illustration and design by Valentin Buckl.



# **Eyes in Motion**

## **A New Traffic Sensing Paradigm for Pedestrians and Cyclists**

**Dissertation**

for the purpose of obtaining the degree of doctor  
at Delft University of Technology  
by the authority of the Rector Magnificus, Prof. dr. ir. T.H.J.J. van der Hagen  
chair of the Board for Doctorates  
to be defended publicly on  
Thursday 1st, May 2025 at 10:00 o'clock

by

**Alphonse Antoine VIAL**

This dissertation has been approved by the promotor.

Composition of the doctoral committee:

Rector Magnificus

Prof.dr.ir. S.P. Hoogendoorn

Prof.dr.ir. B. van Arem

Dr.ir. W. Daamen

Chairperson

Delft University of Technology, promotor

Delft University of Technology, promotor

Delft University of Technology, promotor

Independent members:

Prof.dr. M. Menendez

Dr. J. Kim

Prof.dr. M.J. van den Hoven

Prof.dr. N. Geroliminis

New York University Abu Dhabi, United Arab Emirates

The University of Queensland, Australia

Delft University of Technology

École Polytechnique Fédérale de Lausanne, Switzerland

Prof.dr.ir. J.W.C. van Lint

Delft University of Technology, reserve member

**TRAIL Thesis Series no. T2025/11, The Netherlands Research School TRAIL**

TRAIL

P.O. BOX 5017

2600 GA Delft

The Netherlands

E-mail: [info@rsTRAIL.nl](mailto:info@rsTRAIL.nl)

ISBN: 978-90-5584-366-4

Copyright © 2025 by Alphonse Antoine VIAL

All rights reserved. No part of the material protected by this copyright notice may be reproduced or utilized in any form or by any means, electronic or mechanical, including photocopying, recording or by any information storage and retrieval system, without written permission of the author.

Printed in the Netherlands

*Holy ravioli !*



# Acknowledgements

This research journey has been a winding path through a forest of infinite possibilities, where each step leads to new questions waiting to be explored. Along the way, I was fortunate to have people who guided me, warned me of hidden pitfalls, and, most importantly, believed in me. At times, I felt lost, unsure of my direction, and I can only hope that I did not test their trust too much. Yet, their support kept me moving forward. It is thanks to these people — those I will acknowledge in the paragraphs below — that I found inspiration, learned to see the world differently, and grew both in my work and as a person. Their influence has had a lasting impact on me, and for that, I am deeply grateful.

First and foremost, I would like to express my deepest gratitude to my supervisors, Serge Hoogendoorn, Bart van Arem, and Winnie Daamen. Their support for my curiosity, creativity, and interdisciplinary approach has been invaluable. They created an environment where I felt safe and supported. Serge, your continuous creativity and scientific excellence have been truly inspiring, but above all, I admire your kindness. Thank you for your enthusiasm, for accepting my yellowness, and for the trust you placed in me throughout the years — you have been my mentor in the fullest sense. Bart, you are a role model for academic mentorship and a peaceful guide along this winding path. I deeply appreciate your genuine interest in both my research and me as a person, your ever-open office door, and your encouragement that kept me going when I had doubts. Winnie, your sharp critical eye have pushed me to become a more precise and thoughtful researcher. You taught me to see the details hidden in this forest of knowledge, and I am grateful for your openness in letting me explore perspectives from other fields. Many of the things I learned from you continue to shape my work today, and I often find myself trying to replicate them in my own approach.

Throughout this scientific journey, I've had the privilege of immersing myself in various research fields and communities. I am grateful to those who eagerly embraced new ideas, shared their time and trust, and generously passed on their knowledge. Collaborating across diverse backgrounds, cultures, and expertise is what I love most about research. I feel lucky to have worked with such brilliant and fun-loving people, learning from them and pushing boundaries together. Gustaf Hendeby, you exemplify a model of academic collaboration, while brought me along in the quest to discover Linköping's taverns. I hold great admiration for your interpretation of 'collaboration' working without rigid hierarchies. Seda Gürses, you have profoundly shaped my world-view and development. You showed me how to approach essential work with diligence and the highest standards of precision, opening my eyes to privacy and computational infrastructures in ways I hadn't anticipated. Though things didn't unfold as expected, I've learned some of my most valuable lessons from you. You helped me articulate my thoughts in a new language — one that still shapes my perspective today. Maria Salomons, you have been a reliable collaborator in a twisted project. I am thankful for your continuous enthusiasm

on cycling research and practice. The world needs more cycling-freaks like you. Aaron Yi Ding, thank you for introducing me to your world, sharing ideas, concepts and applications of the distributed computing paradigm. I respect you, your high quality research and would have loved to collaborate more.

I am grateful to the brilliant mentors and collaborators who guided me in ways I never expected. Thank you Jeroen van den Hoven, Carmela Troncoso, Dirk Helbing and Sascha Hoogendoorn-Lanser. I felt that every interaction, no matter the scale, has helped me grow. I truly hope our paths cross again in the future.

The *Allegro* group has been a rich source of support, inspiration and friendship. In particular my fellow *Isotopes-436* including Giulia, Danique, and Tim, as well as, Flo, Vincent (VG7), Marie-Jette, Lara, Alexandra, Yan, Martijn, Lucia, Dorine, Victor, and Yufei. There are countless memories I could mention, but I'll keep them in my heart rather than on a white piece of paper. Thank you to TU Delft Transport and Planning members who shared and built ideas along the way. Special thanks to Maria, Niharika, Tin, Pablo, Yihong, Ding, Xavi, Bahman, Konstanze, Peyman, Hans, Simeon, Oded, Andreas, Gonçalo, Panchamy, Niels, Meng, Peter, Edwin, Priscilla, Dehlaila, and Moreen. The time with you was colourful and abounding with beautiful anecdotes on campus, at events, cafes, restaurants, parties, and on the football field. Of course, I must mention the *DCF Inter* team — seriously, what an honour playing with you! You guys really got me back into the game and reignited that competitive spark.

And finally, it is the deep love, care, and support from my family and dearest friends that nourish and sustain all my actions. I am lucky to have you — I love you.

Alphonse  
Rotterdam, the Netherlands

# Contents

|  |            |
|--|------------|
| <b>Acknowledgements</b>  | <b>vii</b> |
| <b>1 Introduction</b>  | <b>1</b>   |
| 1.1 Active mode data for research and practice . . . . .           | 1          |
| 1.2 Knowledge disparity . . . . .                                  | 2          |
| 1.3 Data collection and processing . . . . .                       | 3          |
| 1.3.1 Contemporary data collection techniques . . . . .            | 4          |
| 1.3.2 Autonomous dreams and a new sensing paradigm . . . . .       | 4          |
| 1.3.3 The unexplored dimensions . . . . .                          | 5          |
| 1.4 Research objective and questions . . . . .                     | 6          |
| 1.5 Research approach . . . . .                                    | 8          |
| 1.6 Contributions . . . . .  | 10         |
| 1.6.1 Contributions to science . . . . .                           | 10         |
| 1.6.2 Contributions to society . . . . .                           | 11         |
| 1.7 Thesis outline . . . . .                                       | 12         |
| <b>2 Towards new system designs for active mode mobile sensing</b> | <b>15</b>  |
| 2.1 Introduction . . . . .   | 16         |
| 2.2 Problem formulation . . . . .                                  | 17         |
| 2.2.1 Different data for different active mode research . . . . .  | 17         |
| 2.2.2 Active mode data collection methods . . . . .                | 18         |
| 2.3 Fundamentals . . . . .   | 19         |
| 2.3.1 Illustrative example . . . . .                               | 19         |
| 2.3.2 Requirements and performance metrics . . . . .               | 20         |
| 2.4 Design of an active mode sensing network . . . . .             | 22         |
| 2.4.1 Sensor network characteristics . . . . .                     | 23         |
| 2.4.2 Functional architecture . . . . .                            | 25         |
| 2.5 Active mode sensing applications with urban vehicles . . . . . | 29         |

|          |   |           |
|----------|---|-----------|
| 2.5.1    | Sensor configuration . . . . .  | 29        |
| 2.5.2    | Active mode sensing scenarios . . . . .                                 | 31        |
| 2.5.3    | External operating conditions . . . . .                                 | 33        |
| 2.6      | Conclusion . . . . .  | 34        |
| <b>3</b> | <b>Network-constrained tracking of cyclists and pedestrians</b>         | <b>35</b> |
| 3.1      | Introduction . . . . .  | 36        |
| 3.2      | Background . . . . .  | 38        |
| 3.2.1    | Bayesian filtering . . . . .  | 38        |
| 3.2.2    | Multiple target tracking . . . . .                                      | 39        |
| 3.3      | Problem formulation . . . . .   | 40        |
| 3.3.1    | The network constraint: road network . . . . .                          | 40        |
| 3.3.2    | Network-constrained state-space model . . . . .                         | 41        |
| 3.3.3    | False observations and new targets . . . . .                            | 42        |
| 3.4      | Network-constrained multiple hypothesis tracking . . . . .              | 42        |
| 3.4.1    | Extended Kalman filter . . . . .  | 43        |
| 3.4.2    | Track hypotheses and scores . . . . .                                   | 45        |
| 3.4.3    | Global hypotheses . . . . .   | 47        |
| 3.4.4    | Complexity saving methods . . . . .                                     | 48        |
| 3.5      | Simulation study . . . . .  | 49        |
| 3.5.1    | Simulation setup . . . . .  | 50        |
| 3.5.2    | Variable negative information . . . . .                                 | 52        |
| 3.5.3    | Results . . . . .   | 52        |
| 3.6      | Outlook . . . . .   | 58        |
| 3.6.1    | Motion and target dynamics . . . . .                                    | 58        |
| 3.6.2    | Sensing infrastructures and measurements . . . . .                      | 58        |
| 3.6.3    | Privacy considerations . . . . .  | 59        |
| 3.7      | Conclusion . . . . .  | 59        |
| <b>4</b> | <b>A novel data source for advanced traffic management and control</b>  | <b>61</b> |
| 4.1      | Introduction . . . . .  | 62        |
| 4.2      | Approach . . . . .  | 64        |
| 4.3      | Inferring the number and states of indistinguishable cyclists . . . . . | 65        |
| 4.3.1    | Road network-constrained MTT . . . . .                                  | 65        |



|          |   |           |
|----------|---|-----------|
| 4.3.2    | State space models . . . . .  | 66        |
| 4.3.3    | State prediction and update . . . . .                                       | 68        |
| 4.3.4    | Applicability and assumptions in a traffic signal control setting . . . . . | 68        |
| 4.4      | Arrival time and priority request . . . . .                                 | 69        |
| 4.4.1    | Arrival time calculation . . . . .  | 69        |
| 4.4.2    | Conditional priority request . . . . .                                      | 69        |
| 4.5      | Reducing lost time by phase truncation . . . . .                            | 70        |
| 4.6      | Experimental design and simulation setup . . . . .                          | 71        |
| 4.6.1    | Simulation scenarios . . . . .  | 71        |
| 4.6.2    | Case study . . . . .  | 72        |
| 4.6.3    | Network representation . . . . .  | 72        |
| 4.6.4    | Field of view and occlusion . . . . .                                       | 73        |
| 4.6.5    | Implementation details . . . . .  | 74        |
| 4.7      | Results . . . . .   | 74        |
| 4.7.1    | Q1: VAT vs. baseline control scenarios . . . . .                            | 75        |
| 4.7.2    | Q2: bicycle green actuation, bicycle and car green duration and green use   | 77        |
| 4.7.3    | Q3: impact of CAV penetration rate and field of view . . . . .              | 79        |
| 4.7.4    | Q4: what is the value of observing the absence of cyclists? . . . . .       | 80        |
| 4.8      | Conclusion . . . . .  | 80        |
| 4.8.1    | Future work . . . . .   | 81        |
| <b>5</b> | <b>Privacy challenges introduced by intelligent autonomous systems</b>      | <b>83</b> |
| 5.1      | Introduction . . . . .  | 84        |
| 5.2      | Comparative analysis with established systems . . . . .                     | 87        |
| 5.2.1    | Location-based systems . . . . .  | 88        |
| 5.2.2    | Crowdsourcing systems . . . . .   | 89        |
| 5.2.3    | Intelligent autonomous systems . . . . .                                    | 90        |
| 5.3      | What's next . . . . .   | 95        |
| <b>6</b> | <b>Conclusions, implications and recommendations</b>                        | <b>97</b> |
| 6.1      | Main findings . . . . .   | 97        |
| 6.2      | Overall conclusions . . . . .   | 99        |
| 6.3      | Implications for practice . . . . .   | 100       |
| 6.4      | Recommendations for future research . . . . .                               | 102       |

|   |            |
|---|------------|
| <b>Bibliography</b>                     | <b>105</b> |
| <b>Summary</b>                          | <b>119</b> |
| <b>Samenvatting (Summary in Dutch)</b>  | <b>123</b> |
| <b>About the author</b>                 | <b>127</b> |
| <b>TRAIL Thesis Series publications</b> | <b>129</b> |

# Chapter 1

## Introduction

About fifty years ago, the marketing industry was guided by the principle *You are where you live* that the location of a person is a defining factor of their identity (Allen, 1982; Klein, 1989). This principle led to the development of the initial geodemographic systems. By merging public data with private computing resources, marketers could segment areas of interest into focused zones, enabling a more precise profiling of individuals, households, and neighbourhoods based on specified categories (Phillips & Curry, 2002). These practices can be seen as precursors to current location data collection methods, where advances in ubiquitous sensing and communication technologies have added a digital dimension to urban environments, going beyond the mere capture of geographic features.

Modern digital technologies enable the collection of a much wider range of data than traditional zip code targeting, a significant leap in the scale, precision, and potential applications of such data. In 2017, Sidewalk Labs, a Google-affiliated company, envisioned a futuristic city when it promised to make Toronto’s Waterfront district *the most measurable community in the world* (Bliss, 2018), gathering extensive data on the environment, transportation, energy and infrastructure. The pursuit of increasingly large, sophisticated, and potentially invasive forms of data collection is often enough driven by hopes of addressing pressing issues (*e.g.*, air pollution, traffic congestion, or even public safety) with fancy technology and shiny new datasets — overlooking the importance of a more nuanced understanding of the impact deployed digital technology has on humans, their social dynamics, and the built environment. The termination of Sidewalk Labs’ project (Jacobs, 2022) highlights the critical balance between the potentials and limitations, risks, and harms that can stem from extensive data collection and analysis technologies in the lived environment.

### 1.1 Active mode data for research and practice

While aiming for smarter, safer, and more efficient urban environments, harnessing the power of sensor measurements and computed predictions about human mobility has become a central aspect of modern traffic monitoring, control, and management. Digital footprints of humans in urban space have been useful in designing and planning transportation and infrastructure projects that are often expensive, contributing to safe, sustainable, and equitable transport. The increasing availability of data provides researchers and practitioners with a valuable resource in

shaping more liveable and sustainable cities, where walking and cycling are viable alternatives to car use, although data alone does not guarantee success.

Walking and cycling are basic modes of transport, and principal elements in the set of mobility systems. In urban transportation, pedestrians and cyclists contribute a unique layer of complexity. A major challenge in contemporary traffic and transportation theory, however, is that there is still much to learn about how, when, and where people walk and cycle. The movements of active modes (*i.e.*, walking and cycling), are complex due to the many degrees of freedom, unlike network-constrained car traffic. Theories have allowed scholars to effectively analyse and predict the demand and the behaviour of pedestrians and cyclists at different levels (Hoogendoorn & Bovy, 2004; Gavriilidou et al., 2019a), while various modelling techniques have been introduced to examine these different behavioural levels, both at a micro- and macroscopic scale (PTV, 2016; Duives et al., 2019; Tan et al., 2024).

From the perspective of transportation researchers and practitioners, data about pedestrians and cyclists are not only essential for observing their behaviour in urban environments, but also for understanding the impact they have on traffic flow and road safety. Data are crucial for developing predictive models that can help, for instance, better design transportation infrastructure, or improve equity and access to safe and efficient pedestrian and cycling infrastructure. When available, pedestrian and cyclist data can be utilized in different forms and types (Feng et al., 2021), ranging from macro-level analyses of the entire transportation network to micro-level investigations of individual links or cross-sections. This enables researchers to validate their models, and allow practitioners and policymakers to use advanced tools and make well-informed choices that align with the latest developments of their respective fields.

## 1.2 Knowledge disparity

In the last decade, there has been growing interest in studying walking and cycling as an essential mode of transportation. This interest, combined with augmented awareness of goals, strategies, and management has led to more research focus and real-life implementations, supported by increased funding. Despite the acknowledged importance of the field (Sun & Yin, 2017), however, researchers and practitioners still face a relative lack of empirical approaches and advanced tools to explain and predict urban walking and cycling dynamics. The long-standing historical focus on motorised vehicular modes has constrained the availability of high-quality datasets and techniques to collect, process and analyse these data, impeding the empirical study of active mode dynamics.

With the strong industry-backed funding of car research and governmental focus on motorised vehicle oriented roadway infrastructure from the past, it is only logical that this field is still in its infancy. The lack of investment and attention has resulted in a relative scarcity of sensors, methods, data, and information in the active mode field. Notably, long-lasting effects of continuous car-centric investments are visible in the pervasive traffic surveillance infrastructure generating data 24/7; non-existent at this scale for active mode traffic, while the availability of

such data is limited in most parts of the world<sup>1</sup>. Ultimately, the scarcity of pedestrian and cyclist data assumably induced a lack of micro- and macroscopic models, as it plays a crucial role in calibrating and validating new models, which in turn limited their availability. Furthermore, due to the scarcity of data, methods for data collection and empirical analysis remain limited, inducing a restricted number of specific types of empirical active mode studies and insights.

Undoubtedly, pedestrians and cyclists exhibit intricate movements, characterized by their multiple degrees of freedom and their complex behaviours, inherent to their continuous social interactions within diverse contexts. This uniqueness challenges the collection of data compared with traditional methods used for vehicles like cars. Moreover, the smaller size and reduced iron content of individuals (in contrast to cars) further complicates accurate monitoring and analysis. It should also be noted that experiments in controlled environments often have limited scope and struggle to represent complex urban scenarios, while pedestrians and cyclists may have a different expectation of privacy than car drivers in public spaces. This poses challenges and may induce more sophisticated sensor settings and methods able to capture and process the specific data requirements from the field. In this context, the fusion of multiple data sources and combined techniques to convert it into a form that can be interpreted and analysed by humans<sup>2</sup> has great potential for improvement. Drawing from other fields such as computer vision, signal processing, or simulation, innovative methods could emerge in the future to enhance this process even further in the field of active modes.

In the broadest sense, the ultimate promise of data collection and analysis in active mode transportation, is to gain a better understanding of pedestrian and cycling behaviour, identify areas of improvement, and develop evidence-based policies and strategies that prioritize safety and efficiency. This shall not only benefit the individuals who rely on these modes of transportation but also contribute to the overall health and well-being of our communities. Both the lack of data and robust processing techniques, however, may be factors contributing to the relatively restricted scientific and practical advancements in understanding and developing active mode traffic to the same levels as car traffic. This constraint also manifests in the mostly timid implementation of these findings in cities globally.

## 1.3 Data collection and processing

As we have described earlier, it is likely that the scarcity of data and processing techniques for studying pedestrian and cyclist movements has impacted the availability of new theories, models and tools to explain and predict urban walking and cycling dynamics in urban contexts. However, a new wave of technological advancements, driven by industry and research communities from different fields, is now emerging. The developments are likely to have a profound effect on the landscape, in particular on the collection and analysis of data using many different methodologies. Processing this data enables measuring physical quantities or classes such as presence, pressure, acceleration, location, shape, object type, and so on. These may be used to

---

<sup>1</sup>Note the role of geographic disparity of data in AI (Sambasivan et al., 2021), also reflecting in urban development where areas with more extensive data collection infrastructure, may have an advantage in urban planning and development (Malaker & Meng, 2024).

<sup>2</sup>When frequently the problem is much earlier in the process as most people do not know what to look for or are unable to phrase appropriate questions.

estimate spatiotemporal densities, velocities, or flows of active modes; providing information to an end-user or as input to another system.

### 1.3.1 Contemporary data collection techniques

In the past, traditional methods such as manual surveys and observational studies have been complemented by more advanced digital technologies. The integration of sensors, video analytics, and GPS tracking systems has enabled researchers and practitioners to experiment with, or deploy data collection systems to gather, sometimes in real-time, sometimes high-resolution data on pedestrians and cyclists. A brief look at the research output within the transportation community underscores the increasing number of scientific contributions enabled by an array of different data collection methods in the past few years. These include, field studies employing GPS-equipped pedestrians and bicycles (Daamen et al., 2017; Strauss & Miranda-Moreno, 2017), controlled experiments with static vision sensors (Gavriilidou et al., 2019b), the analysis of (big) social media data (Yang et al., 2019; Gong et al., 2020), incorporation of Wi-Fi and Bluetooth technologies (Versichele et al., 2014; van Oijen et al., 2020), the integration of more involved Radar and LiDAR sensors (Lee et al., 2020), or more recently, the utilization of virtual reality (VR) (Feng et al., 2024). Each techniques offers a unique perspective on movement patterns, interactions with urban space and infrastructure, and responses to other traffic participants or environmental factors. Despite smartphones and other wearable devices enabling large-scale data collection (*e.g.*, capturing population movements on a macro-scale (Schläpfer et al., 2021)), the proprietary nature of most mobile applications makes such data difficult to access for researchers in the public sector.

### 1.3.2 Autonomous dreams and a new sensing paradigm

The digital layer capturing pedestrian and cyclist spatial and temporal characteristics with contemporary sensing techniques may soon be augmented by a new sensing paradigm, which could offer the availability and accessibility of data that seems so essential to foster innovation and sustainable impact to take place. The rise of distributed computing and connectivity infrastructure, but especially, the growing sensing and processing needs of intelligent autonomous systems (IAS), lead to an exponential increase in the generation of captured human features — such as their movement patterns, shapes, and contextual information — at different temporal and spatial scales. Connected autonomous vehicles (CAV) are an instance of intelligent autonomous systems that can process and analyse data from several sensors to make decisions and perform tasks independently. The promise of these new systems has captured the imagination of the many, and the hope extends beyond mere technological innovation. Some foresee a future where safety is drastically increased through a decrease of road accidents (Hawkins, 2023). Others find inspiration in the prospect of diminished urban traffic (Lang et al., 2020), or in the development of a more inclusive and accessible transport system (Aquino, 2024).

CAVs or self-driving cars, for instance, pushed by seemingly unstoppable private and political forces, are being developed and used progressively and increasingly. Real-world deployments on roads have rapidly materialized, showcasing the tangible progress in intelligent and autonomous vehicle technology. Companies like Waymo, Cruise, Wayve, and others have taken

significant strides, placing semi-autonomous and autonomous vehicles on public roads for testing and, in some cases, commercial use. They use perception systems that leverage advanced data and compute capabilities of deep learning, producing an enormous amount of sensor data (approximately 750 megabytes per second (Harris, 2022b)). A necessary precondition to deploy such data-hungry systems, is the availability of powerful computational resources to perform calculations at very high speeds, possibly in real-time<sup>3</sup>. This abundant stream of data is utilized as a valuable source of training material for deep learning algorithms to enhance process this data and increase their performance in driving tasks. The integration of 5G/6G connectivity should enable collaborative approaches to analyse vision or other sensor data at the edge/fog, leading to more robust and scalable data ways of identifying and classifying objects, such as pedestrians within the environment, providing a crucial layer of perception for autonomous driving systems.

The motivation for this thesis is driven by the exciting possibilities and inherent limitations and challenges, emerging from this new sensing paradigm, pushed by these technological innovations and autonomous dreams. It should be made clear, however, that this thesis not only applies to autonomous vehicles, but also to more traditional moving<sup>4</sup> sensor platforms (MSP) that do not necessarily operate autonomously, but are equipped with a wide range of sensors. Such human-operated ground or airborne vehicles (*e.g.*, sensor-equipped cars, bicycles, or even humans with bodycams) may also be equipped with advanced sensors and technologies to collect and analyse immense amounts of data about their surrounding environment, yet do not rely on this information to self-operate.

### 1.3.3 The unexplored dimensions

In the context of IAS, it is important to recognize these systems are specifically designed to capture spatiotemporal data of objects in urban environments (and thus also humans) as they require this information to self-operate. Justified by their primary function (*i.e.*, self-operation) they require to capture their environment to maximize the safety of all traffic participants, and thus coordinate massive amounts of data across time and space. Therefore, in a world where pedestrians and cyclists are (still) an integral part of transportation systems, the success of IAS inherently hinges on their ability to capture these individuals. If unable to capture surrounding people despite their presence, that is, finding the optimal balance between detection accuracy, false positives and false negatives, it is effectively useless, rendering the entire technology obsolete.

The generation of pedestrian and cyclist data required to self-operate IAS, not only allows to improve their situational awareness but may also enable a variety of secondary functions and applications. It creates a by-product that may be exploited for other purposes by traffic researchers and practitioners, such as traffic monitoring, control and management. For example, the generated data may be a source of measurements and insights into how and where people move on sidewalks and cycle paths. This information would allow to obtain a view of micro-

<sup>3</sup>The necessary infrastructure deployment may not be as much of an issue as one might think, given the significant investments being made in this area.

<sup>4</sup>Although different research communities may accord nuances in the exact definition of "mobile" and "moving," this work uses these terms interchangeably to refer to systems that can collect data while in motion. This can include humans, vehicles, drones, and all other forms of transportation-based sensing systems.

scopic and macroscopic traffic flows in an observed area, an asset for advanced surveillance applications in traffic operation, control, and management. Inspired by previous work in the context of environmental monitoring with traditional moving sensor platforms (Anjomshoaa et al., 2018; Minet et al., 2018), this thesis speculates on the secondary functions of rising IAS (or traditional moving sensor platforms as dedicated forms of data collection) to emerge as a new source of data for traffic research and practice.

With this in mind, leveraging this technological push to address still unexplored pedestrian and cyclist domains comes with a set of challenges and unknowns when it comes to the potential utilities and societal impacts. In particular in the context of pedestrian and cycling traffic research and practice, where significant gaps emerge when mapping current domain needs to technological advancements emerging from IAS. In theory, present needs and challenges from the active mode community may be addressed with emerging wealth of information; however, how this data may technically be exploited, and how to make information from all these data, remains unclear. Further, existing application scenarios and future use cases must be (re)evaluated, as this augmented data source likely provides a different set of inputs that, with today's data-driven approaches, may affect the output of the systems. The practical use and deployment also questions the societal impact in this (brave new...) sensing paradigm, in particular as we collect data about a-priori not explicitly consenting human beings. From the perspective of the traffic engineering community, this nascent interdisciplinary field opens up a large blank canvas, a unique opportunity to integrate insights from various disciplines that can transform the way we capture, interact, and navigate with our surroundings.

## 1.4 Research objective and questions

This thesis seeks to investigate the potential utility and societal impact of exploiting pedestrian and cyclist data from moving sensors (IAS or MSP) in the context of traffic research and practice. The use of the term *exploitation* is twofold: 1) the development of methodological knowledge and tools to extract traffic information from the data, and 2) under what conditions (*e.g.*, limitations, risks, and harms) it contributes to technological, societal and economic development. To fulfil the main objective of the research, several derived research questions were formulated. In the following, the research questions and the relation between them is detailed:

**What are the technical and functional requirements to a pedestrian and cyclist traffic sensing system that uses moving sensor platforms to collect data?** The new digital layer created from by-products of the emerging sensing paradigm is undertheorized for traffic engineering applications. It is thus essential to better understand the origin of the data and the interplay of related infrastructure. This motivates to first define and articulate both the technical and functional prerequisites essential for a sensing system using a connected network of mobile platforms as sensor nodes to capture pedestrian and cyclist spatiotemporal properties. This first objective establishes a theoretical foundation by analysing key requirements and constraints, providing a comprehensive overview that outlines the specific features and capabilities necessary for the efficient functioning of the mobile sensing system. This includes an exploration of potential obstacles that may impede the seamless operation of the mobile sensing system, such as data accuracy,



resources, and consumption, as well as the integration of other relevant technologies. [Chapter 2]

**What specific adaptations are required for advanced state estimation and data association methods to effectively infer trajectories from data collected by both static and moving network-constrained sensors?** Having identified different types of data and technical requirements in the previous step, it is possible to envision future mobile sensing scenarios, starting with assumptions aiming at using most basic location information possible in a generic road-network context. Presuming the availability of the pedestrian and cyclist measurements, the next step involves exploring methodologies to effectively apply and derive insights from the future data. A second objective of this thesis is thus to conduct a comprehensive review of the current state-of-the-art methodologies related to state estimation and data association techniques employed in inferring trajectories using data obtained from static and moving sensors. A systematic analysis of identified gaps grounds the research and enables to contribute to the advancement of knowledge in the field by proposing a novel methodology that addresses the limitations of existing approaches, ultimately enhancing the applicability of trajectory inference from pedestrian and cyclist data collected by network-constrained moving sensors. [Chapter 3]

**What potential advancements in traffic control and safety can be achieved by integrating data derived from moving sensors such as connected autonomous vehicles?** By understanding the functional challenges and requirements of the sensing system, and developing a method to use the data adequately, the tangible benefits and practical applications of integrating data derived from moving sensors into future traffic surveillance systems can be explored. In light of the increasing attention from the traffic community towards reducing delays for cyclists at signalized intersections, this investigation seeks to demonstrate how data from dynamic sensing infrastructure can provide valuable insights and help optimize traffic control and management strategies to improve road safety and contribute to the development of intelligent transportation systems. The practical implications will be investigated with a concrete use case at hand, emphasizing the potential to manage traffic flow in contemporary urban environments with this new type of information. [Chapter 4]

**What are potential privacy concerns affecting pedestrians and cyclists that emerge from traffic engineering use cases utilizing IAS as they enter urban environments?** After exploring the potential benefits and functionalities of this new sensing paradigm, it is crucial to study the interplay and its implications within the realms of political-economical forces, privacy and security, and traffic research. By positioning these inter-related potentials and conflicts within the purview of diverse research communities, this study seeks to highlight the multifaceted dynamics at play. This study investigates how mobile sensing technologies intersect with political and economic structures, exploring their impact on societal frameworks, and investigates the balance between the potential utility of mobile sensing and the associated concerns surrounding privacy. Ultimately, the objective is to understand how IAS infrastructure is different from other digital infrastructures, stimulating future research to address shown gaps, and guide future deployments in ways that align with ethical values. [Chapter 5]

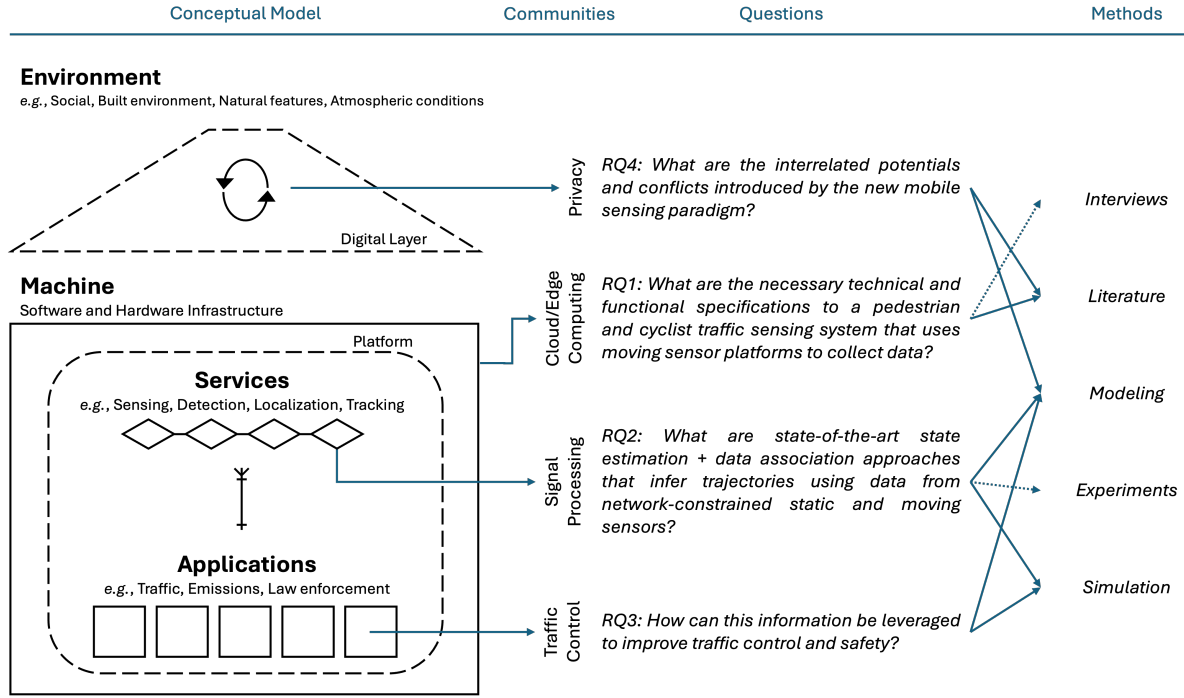


Figure 1.1: Research approach

## 1.5 Research approach

In this section, the research approach illustrated in Figure 1.1 will be described, and grounds on a conceptual framework proposed by Jackson & Zave (1995). Building on their work, we define a *system* as the *machine* together with its users and the whole *environment*. We apply parts of this influential conceptual framework both to the context of IAS and non-autonomous MSP, to clarify the system boundaries, the requirements, specifications, as well as the elements that the system of study interacts with. By analysing the key elements and relationships within this framework, we can gain valuable insights into the underlying mechanics, systematically identify areas of research, and delineate the different elements of this thesis.

In the context of this thesis the term *machine* describes the hardware and software apparatus including the material infrastructure as well as data platforms built on top of these infrastructures. We consider *sensing, positioning, processing and communication infrastructures* to be the material infrastructure needed [for IAS or MSP] to generating measurements about the machine’s internal and external state, retrieving the machine’s position, leveraged for data collection and processing, and sharing information across the system. The environment, in contrast, refers to the portion of the real world relevant to the operation of the machine, e.g., traffic participants such as pedestrians, cyclists, and drivers, as well as the built environment comprising roads, buildings, and other infrastructure. Additionally, it includes natural features like vegetation, atmospheric conditions, lighting, and other factors. Furthermore, social and economic dimensions may also be included within this broader definition of the environment. At a given time, a sensor generates a *capture* of the machine’s internal or external state (e.g., to obtain *measurements* about distance, speed and colour of an external object). We refer to a *scan* as the set of all captures received by a sensor at a given time, while the machine’s internal or external state may be captured by multiple sensors at the same time instant.

The infrastructures are interconnected in different layers and represent how the physical compute is organized for the functioning of the machine in the environment. A *platform*, on the other hand, can be described as the operating system (OS) like layer that enables and is a gateway to services that make use of the different material infrastructures. This means that, platformization enables various *services* that leverage resources provided by the different infrastructures, *e.g.*, data collection, localization, detection, tracking, or control services. The computation initiated by services can be organized in different ways, *e.g.* on-device or centralised server, while it dynamically leverages sensing, positioning, processing, and communication capabilities of cloud, edge, and mobile devices to form a digital layer that can be used for different *applications* such as, navigation, gaming, advertising, emissions monitoring, parking fees, or traffic monitoring and control applications. Infrastructure and platform are intertwined and we distinguish both in the ways they are controlled<sup>5</sup>, that is, they can be owned and operated by the same or different parties<sup>6</sup>.

To comprehensively investigate and gain a holistic understanding of this emerging field, an interdisciplinary research approach was adopted, integrating transport research with other relevant scientific communities. To address Research Question 1 effectively, relevant data and perspectives (which are not immediately evident or accessible within existing publications) were compiled to provide a comprehensive overview of the topic of study. Consequently, desk research entailed a literature review as well as interviews with domain experts, aiming to systematically compile and synthesize current knowledge across related disciplines, while characterising design constraints and functional requirements to the system of study<sup>7</sup>. Recognizing the limitations of state-estimation techniques applied in traditional traffic contexts, we expanded the scope by studying their foundational aspects — signal processing and tracking theories — thereby constructing a robust framework for fusing pedestrian and cyclist observations to address Research Question 2 through experiments and simulation validation. To answer Research Question 3, a traffic control scenario was defined, while modelling and simulation were used to generate scenario-specific data at scale. Due to the lack of adequate real-world data, simulation was used when necessary to address the research questions 2, 3 (and 4, however not included in this thesis). Compensating for the emerging nature of this sensing paradigm, filling gaps within location privacy research, and acknowledging the absence of standards, addressing Research Question 4 involved examining the broader socio-technical impact through a review of academic literature, but also including non-scientific magazine and journal articles and opinion pieces, to support the conceptual modelling of the system.

---

<sup>5</sup>For example, entities providing servers, mobile devices, sensors and processors, as well as application programming interfaces (APIs) as the basis for all software development and applications (Apple, 2020). A reflection on consequences of power asymmetry and control over users, developers and increasingly, public institutions is, however, outside the scope of this chapter.

<sup>6</sup>We could argue, based on (Plantin, 2018), that a platform with the data it makes available is a knowledge infrastructure. A discussion about the control and consequences is, however, outside the scope of this thesis.

<sup>7</sup>Jackson & Zave (1995) refer to a *control system* as “a machine that interacts with its environment to bring about or maintain relationships in that environment”. Thereby, requirements state these desired relationships, and are concerned entirely with the environment, where the effects and benefits of the machine will be felt and assessed. This definition aligns with our approach and understanding of functional challenges and requirements

## 1.6 Contributions

The overarching contribution of this thesis is to enhance knowledge on extracting actionable pedestrians and cyclists information from data collected by moving sensors (IAS or MSP), as well as its related societal impacts. Findings from this thesis contribute to science by developing methodological knowledge on information extraction and highlighting previously overlooked implications on privacy and beyond. This thesis has implications for practice, especially in areas facing increasingly high levels of sensor deployments, by providing insights for policy development, traffic management, and responsible innovation.

### 1.6.1 Contributions to science

All studies contain one or multiple contributions to science. These are presented below in the order-of-appearance in this thesis.

#### **Insights into requirements, challenges and design constraints for systems using pedestrian and cyclists data from mobile sensing platforms (Chapter 2)**

The generic characterization of the system, yet described from the perspective of traffic research and practice, could contribute to the scientific understanding of novel forms of large-scale data collection in various fields such as engineering, computer science, environmental monitoring, or even epidemiology and economics. The analysis of functional requirements could provide researchers with insights into the minimum set of features needed for such a system to be effective (which may vary based on factors such as data collection objectives, environmental conditions, regulatory frameworks, etc.), identify gaps between data needs and computational infrastructure, and contribute to methodological advancements in data collection and integration. The identification of functional challenges and design constraints could inform future research on how to balance these factors in system design. The theoretical design integrates the core components of the system as a starting point for relevant communities to work on challenges related to reliability, scalability, ethical, privacy, and security issues inherent to this design, while exploring potentials such as of context-aware methods in data collection.

#### **Methodology for network-bound targets in multi-target tracking problem (Chapter 3)**

The proposed framework enables the reconstruction of trajectories from an unknown and varying number of pedestrians and cyclists across larger networks, also given false measurements and clutter. Building upon this framework, basic location information from any moving objects or subjects can be fused together to infer individual trajectories. The inclusion of the network-constraint may be used as building block when extending or developing new state-estimation and data-association algorithms, since it has been shown to allow for more efficient predictions over extended periods of time and a simplification in the measurement association process. With the flexibility for multiple sensor modalities and settings, the proposed framework offers the possibility to include the environment more comprehensively and could be applied to various fields.

#### **Control method to inform, prioritize and optimize an intelligent traffic signal controller (Chapter 4)**

The impact of integrating cyclist data generated from CAVs on traffic signal control (depending on specific parameters, *e.g.*, settings of the CAV, penetration rate of the CAV, occlusion)

is explored with a proposed vehicle-actuated controller approach, incorporating tracking capabilities. The effect of cyclist demand on the delay of cyclists and cars evaluated. Using a straightforward approach, which would not require significant upgrades to existing infrastructure, introducing CAVs at low penetration rates can already help reduce cyclist delays. Furthermore, observing the absence of cyclists appears to be valuable information to the controller, as it enables phase truncation to minimize lost time. These findings and insights have implications for the scientific community, with potential applications in reducing traffic congestion, improving traffic flow, and enhancing road safety. Combining novel data collection approaches with adapted control methods, may also inspire future research to dynamically prioritize specific transport modes.

### **Foundation for investigating location privacy implications for non-users that stem from the infrastructural shift introduced by IAS (Chapter 5)**

Using the conceptual framework presented in Chapter 1.5 helped define appropriate terminology and better relate interactions within the system. The presented comparative analysis of IAS and other established digital systems highlights how the relationship between the machine and the environment is different for IAS than from other existing systems. Identifying this area overlooked in scientific discourse, which implies location privacy implication for active modes, can serve as an initial foundation for interdisciplinary work across scientific communities (*e.g.*, transportation, privacy, security, or law). The findings particularly support future research, and motivate towards the development of more comprehensive privacy analyses, including the evaluation of existing privacy frameworks and models.

## **1.6.2 Contributions to society**

This dissertation contains multiple societal contributions. These are presented below in an order that is not specific.

### **Opportunities for transforming transportation systems**

Using data generated by IAS or MSP could help improve traffic safety, efficiency, or even accessibility within the network. Furthermore, incorporating additional data such as traffic signs, green spaces, or mobility/public transit information may help identify areas where infrastructure improvements are needed. For example, spotlighting hidden commuting pathways to build new lanes for faster cyclists. Building on the systematic analyses of data requirements and technical functionalities, private and public institutions can design and develop technological sensing systems, evaluate necessary infrastructure investments and research efforts. Further, researchers and practitioners can build proof-of-concepts following the principal components of this thesis, to optimize their transportation systems, offering more adequate mobility solutions, *e.g.*, better understanding demand (also predicting).

### **Guidebook and protocols**

Furthermore, this thesis contributes by advancing a set of requirements and design challenges that are elementary when using the data in the context of traffic engineering. It can help practitioners understand the potential of using moving sensors platforms to capture pedestrian and cyclist information. It enables getting a sense of the data collection, at different spatial and temporal levels, given technological possibilities and urban constraints. This is valuable when assessing the potential deployment of such systems in urban areas. In fact, self-driving car com-

panies already enter partnerships with cities in which self-driving cars are unleashed on public roads, while in return, they may be willing to share certain data (Grauerholz, 2017; Hawkins, 2019).

### **Advancing traffic monitoring, control and management**

A more methodological contribution can be found in the framework proposed to an unknown and varying number of pedestrians and cyclists in a road-network. The presented network-constrained multi-target tracking methodology, tracks using information from moving (and can be easily extended to static) sensors. The flexibility of the framework makes it applicable in real-world applications, *e.g.*, information from different sensors can be used if they provide individual position on the road and the general map, and different motion or measurement models can be applied. The proposed methodology serves practitioners a foundation for initiating experiments of advanced traffic applications (*e.g.*, intersection control, monitoring) in centralized data processing settings. For example, by leveraging state-of-the-art traffic signal systems (*e.g.*, those currently in use in the Netherlands) this approach could be incorporated into current systems to enhance capabilities by processing data from moving sensors.

### **Surveillance risks and harms in IAS**

The analysis and formalization of the new form of continuous capture of the public environment including non-users, where different services are foreseen to power a set of potential applications, led to the identification of an infrastructural shift from current computational infrastructures. This provides strong signals to public and private stakeholder to consider a set of novel potential risks and harms led by the deployment of such systems. By signalling the potential expansion of surveillance infrastructures, illustrated with most recent real-world deployments, actors of society are provided with a solid base to reflect and discuss on the future they want to shape.

### **Intuitions on power dynamics**

This dissertation makes contributions to the practice of different actors of the public and private sector, which are part of the IAS paradigm by interests that may be adverse to each other and vie to favour their particular interests and activities. This work is suggestive of using IAS as a source of pedestrian and cyclist traffic data in many different ways; *e.g.*, for designing and planning infrastructure projects (*e.g.*, cycling infrastructure, or access to mobility), crafting new policies and action plans (*e.g.*, spatial community knowledge, or target voters), or strengthening commercial position on the market (*e.g.*, develop new products and services). It is unclear, however, who will ultimately own and manage essential infrastructures and data platforms. While not being its primary focus of this thesis, it initiates a critical reflection into various components the system across its chapters. By unravelling elements to subjective power dynamics — echoes of existing concentration of power within technology companies — this thesis provides policy-makers and other relevant stakeholders with insights that may contribute to a nuanced discourse on governance and oversight necessary for IAS to operate in the public interest.

## **1.7 Thesis outline**

In the following an overview of the chapters of this thesis is presented. Four separate chapters (2-5) address the research questions posed in Section 1.4. These chapters constitute the body

of this thesis and contain articles that were written during the PhD, with the author being the first author. At the beginning of each chapter, it is indicated where the article was published or whether it is under review. The published articles are not modified in any way to avoid potential citation conflicts with the articles found in the journals.

After an introductory chapter, this thesis delves into the various data and compute dimensions that make up the overall sensing system. The reader is provided with an overview of the different components, necessary to capture spatiotemporal properties of pedestrians and cyclists from moving sensor platforms. Concretely, Chapter 2 elaborates on a variety of requirements, along with functional challenges, and outlines the research to be performed with the generated data.

After presenting this overview that can be interpreted as a theoretical design, a more methodological part is provided. Chapter 3 assumes that the data described in Chapter 2 is at our disposal and provides a methodology on how to extract valuable information from it. Specifically, it proposes a new method for advanced traffic applications, tracking an unknown and varying number of moving targets (*e.g.*, pedestrians or cyclists) constrained by a road network, using mobile (*e.g.*, vehicles) spatially distributed sensor platforms.

Chapter 4 then aims to showcase the potential of this new information in a concrete traffic control use case using the method proposed in Chapter 3. Although there are many possible use cases one could think of, an application is selected that prioritizes cyclists at signalized intersections. Data collected from cyclists using mobile sensor platforms is utilized to provide utility back to the same cyclists the data was gathered from.

When collecting data about individuals, utility comes always at a cost of privacy because the more data is collected, the more it is possible to infer sensitive information about the individual. Chapter 5 represents research that not been published yet. Although preliminary, it provides a good starting point to reflect on and discuss identified location privacy implications for pedestrians and cyclists (in this work referred to as non-users) that stem from the infrastructural shift introduced by IAS entering physical space. This chapter provides the reader with a solid foundation and appropriate vocabulary to better distinguish the surveillance infrastructure in development from existing ones, position it in the privacy as well as traffic community, and motivates for future research in this field by presenting unsolvable conflicts once technology becomes infrastructure.

Finally, the last chapter of this thesis concludes with main findings, scientific limitations and discussion, while looking back at the research questions. Finally, the implications for practice and recommendations for future research are provided.





## Chapter 2

# Towards new system designs for active mode mobile sensing

---

The main motivation for this chapter is the need for theory with focus on capturing spatiotemporal properties of pedestrians and cyclists from types of IAS, which is believed to support a host of new research and application development.

With stationary sensor systems or crowd-sourced data from mobile and wearable devices being part of the current active mode data collection toolbox, this chapter investigates both technical and functional prerequisites and challenges of within this novel sensing paradigm. This chapter thus presents a plausible design grounded in specifications on key components such as sensing, networking, processing, and communication, suited to the distinctive context of active mode research and practice. The generic nature of this chapter helps guide different stakeholders with the transfer of their use cases and utility functions in response to the available infrastructure and the surrounding environment. Ultimately, this chapter theoretically grounds design choices and assumptions made in the next chapters, and supports future mobile sensing use cases needed to further identify data and methodological needs.

This chapter is published as a journal article: Vial A., Daamen W., Ding A.Y., van Arem B., and Hoogendoorn S.P. (2020), AMSense: How Mobile Sensing Platforms Capture Pedestrian/Cyclist Spatiotemporal Properties in Cities. *IEEE Intelligent Transportation Systems Magazine*.

---

## 2.1 Introduction

While urbanisation is competing claims on space in cities, it is essential to understand peoples' movements for the design and management of infrastructure, safety, mobility, as well as for public and private transportation. The new digital layer emerging in cities that includes sensors and pervasive mobile systems can help observe and manage different walking and cycling mobility and movement patterns by gathering large amounts of spatiotemporal data. Data about pedestrians and cyclists — so-called active modes — dynamics in urban environments are essential for different types of spatiotemporal analyses, models and behavioural theories. Nowadays for instance, stationary sensing systems are used to continuously monitor pedestrians and cyclists at fixed locations over time (*e.g.*, camera-systems at intersections). Observations about active modes, however, could also be collected with spatially distributed sensing platforms and shared among each other. Although a dense network of stationary observers could possibly meet the desired objective (*i.e.*, providing spatiotemporal mobility information), such large network deployment may require an excessive number of sensor nodes in order to achieve satisfactory sensing performance, at high infrastructural costs. In addition, a static network is not flexible and would not adapt to unpredictable network dynamics or changes of the physical environment. For instance, events due to sensor failure, coverage holes, and changes in the infrastructure or mobility behaviours that are likely to happen in an urban context.

In the past, a number of studies have used vehicles to monitor the urban environment (*e.g.*, traffic, pollution, road conditions). As the number of sensors in a vehicle has increased by the thrive to so-called intelligent vehicles it evolved from a purely mechanical to a genuine cyber-physical system that continuously streams diverse data in real time. Some of these data are essential to the proper working of a vehicle's components and functionalities, but at the same time the captured surplus of data could be used for other purposes (Massaro et al., 2017). The ever increasing number of sensors in intelligent vehicles (*e.g.*, LiDAR, Radar and vision) enable a wide range of urban monitoring applications, thanks to their ample sensing, storage, processing, and communication (*e.g.*, V2V, V2I) capabilities that have not been utilised for this purpose up till now. In the context of this research, an intelligent vehicle can be understood as mobile sensing platform, capturing pedestrian and cyclist spatiotemporal properties from the number of distinct signals it generates. In a foreseeable future, large groups of connected intelligent vehicles, are expected to be deployed in cities and potentially coordinate their actions through communication networks. The promising nature of these connected mobile sensing platforms enables to carry out tasks, proven to be difficult when performed by a single vehicle, or static sensor.

In this paper, complementing current active mode sensing methods, we propose a novel sensing system, called AMSense, that grounds on connected intelligent vehicles as mobile sensing nodes in a network, to capture pedestrians/cyclists spatiotemporal properties in cities. In this dynamic, multi-sensor approach, real-time data, algorithms, and models are fused to estimate spatiotemporal densities, velocities and flows of active modes using presence, position, and movement data collected by a fleet of mobile sensing platforms. Active mode data shall be extracted, processed, and shared through a mobile sensing network.

**Contributions.** AMSense represents the first pedestrians/cyclists mobile sensing system that innovates in such a fundamental part of active mode research, especially in the provision

of real-world, and real-time data. Mobile sensing platforms that collectively gather active mode information in a network, is a rich application domain with many challenges left to be resolved. We therefore characterise design constraints and requirements, in terms of sensing performance, processing, and control, of such a novel active mode sensing system. Probably the most important contributions of this work revolve around the combination of data, the extensive temporal and spatial scale, and the dynamics of the data collection system. This novel sensing paradigm offers a number of advantages over more traditional methods using stationary sensor systems or more recently available data from mobile or wearable devices, as it reduces effort and cost to collect pedestrian/cyclist data, at an extensive temporal and spatial scale, while providing answers to intrusiveness and scaling effects. This paper presents a theoretical design of such a sensing system.

The remainder of the paper is organised as follows: Section II covers different types of active mode research as well as their required data, and compares current data collection methods. Section III illustrates the proposed work with an urban sensing example, and addresses main requirements to such a novel sensing system. Section IV delivers insight into architecture and functions of our proposed mobile sensing network, while Section V elaborates on how to derive active mode spatiotemporal properties in large urban environments. Section VI concludes this paper and highlights future research directions.

## 2.2 Problem formulation

The use of data is not only crucial for the empirical observation of active mode movement behaviour, but at the same time, data is decisive for the development of models, their calibration and validation. Yet, datasets providing comprehensive pedestrian and cyclist mobility information on road-, lane-, or subject-level, are remarkably rare given the rise of sensors in cities. The need for high-quality datasets that capture pedestrians/cyclists in large urban environments is undisputed. This section primarily aims to clarify that different tracks of active mode research require different types of data. Note that we hereby only focus on pedestrian/cyclist movement data collected in urban environments. This section eventually gives an overview of current real-world data-collection methods.

### 2.2.1 Different data for different active mode research

Pedestrian and cyclist dynamics in cities can be described and predicted at three behavioural levels: strategic, tactical and operational (Hoogendoorn & Bovy, 2004), (Gavriliidou et al., 2019a). Several modelling approaches at three behavioural levels have been proposed in the past, both micro- and macroscopically. A comprehensive overview of the main modelling approaches is described in (Duives et al., 2013) for pedestrian models, and (Twaddle et al., 2014) for cyclist models.

Various approaches that study active mode flows and behaviours require pedestrian/cyclist traffic-related data, providing information about movements in space and time. Different models and behavioural theories require different data; mainly varying in spatial scale, accuracy, and granularity. In this context, data granularity represents the scale, or the level of detail of a

dataset, while data accuracy relies on technical capabilities of the data collection system. Furthermore, the temporal character of the data use, or in other words its value loss over time, is related to the aspect of real-time. In essence, studies at strategic and tactical level usually use movement data that was aggregated up to a certain extent, and thus conceivably requires data at lower granularity and accuracy. At the operational level, microscopic models however require particularly detailed (*i.e.*, high granularity and accuracy) movement data along observed individual trajectories, as they examine, for instance, variations in speed, directions, relative positions, or headways. Densities, speeds and flows can however also be observed as fundamental macroscopic relationships, at higher levels of aggregation.

Pedestrian and cyclist data can thus be used in a variety forms and types, each satisfying a distinct track of active mode research. Data-driven studies vary in spatial scale, either over the entire network, a link, or a cross-section, and require more or less time-sensitive information.

## 2.2.2 Active mode data collection methods

Real-world datasets are required for any empirical model of active mode behaviour and analyses of movement patterns. Humans are considered to be dynamic by nature as they move at varying speeds, have different body postures, and wear a range of widely varying clothing that may mask their body shape. These static and dynamic traits can be captured by different types of sensors when an active mode is present (*e.g.*, shape), or performing an activity (*e.g.*, walking/cycling).

Micro- and macroscopic measurements at greater spatial level, however, have become more and more utilised thanks to technological advances linked with mobile technologies. Yet, the emergence of ubiquitous mobile positioning and tracking devices have enabled the gathering of large amounts of data, in a technically convenient and economically cheap way. The movement of pedestrians/cyclists within a given area can be observed from the digital traces generated by a variety of sensing sources, and requires lower accuracy. For instance, frequently used positioning systems (*e.g.*, GPS), are suited to study mobility patterns and route/activity choice, while the variations in their accuracy, where errors can be caused by satellite orbit or receiver issues, make them inappropriate for microscopic movement behaviour analyses. This technological integration hence offers an opportunity to analyse mobility patterns, across transportation modes, and potentially in real-time. Over the last decade, digital footprints of human mobility patterns have enabled urban planners, computer scientists, sociologists and engineers to better understand dynamics in cities. Past studies have mostly used GPS location updates (Shoval et al., 2011; Wirz et al., 2012), mobile phone records (*i.e.*, CDR and sightings data) (Gonzalez et al., 2008; Calabrese et al., 2010), Bluetooth (Versichele et al., 2014), Wi-Fi (Danalet et al., 2014), and social media data (Cheng et al., 2011; Jurdak et al., 2015). Still, these data require direct or indirect collaboration of the tracked target, and thus has to deal with drawbacks in terms of representativity of the sample. But it is the intrusive nature of these participatory mobile sensing methods, relying on smartphones, wearables and other logging devices, which represent a major disadvantage and potential limitation as for data sparseness in some areas. In addition, while phone call records data (mining of mobile phone data) have proven useful for vehicular traffic analysis, they present significant disadvantages due to limited spatial resolution and restricted data access. The spatial resolution may represent a bigger issue for pedestrian, than for cyclist mobility, as cyclists travel further, and hence more coarse data better describes their behaviour. Thereby, mode choices are more difficult to infer than route choices (Chen et al., 2016).

Bluetooth and Wi-Fi are currently referred to as the most suitable technology to measure active mode traffic in a broad set of traffic conditions: from low traffic volumes to high volume flows in an urban network. They allow tracking of individuals through a network, one of their main advantage being the low cost of measurements (Daamen et al., 2015). Wi-Fi, however, appears to be more generic and thus more suitable. At the same time, owing to the pervasive coverage, low latency and high bandwidth, the fifth generation (5G) of mobile networks are becoming another promising solution for urban scale mobility measurement (Wang et al., 2014; Ding & Janssen, 2018).

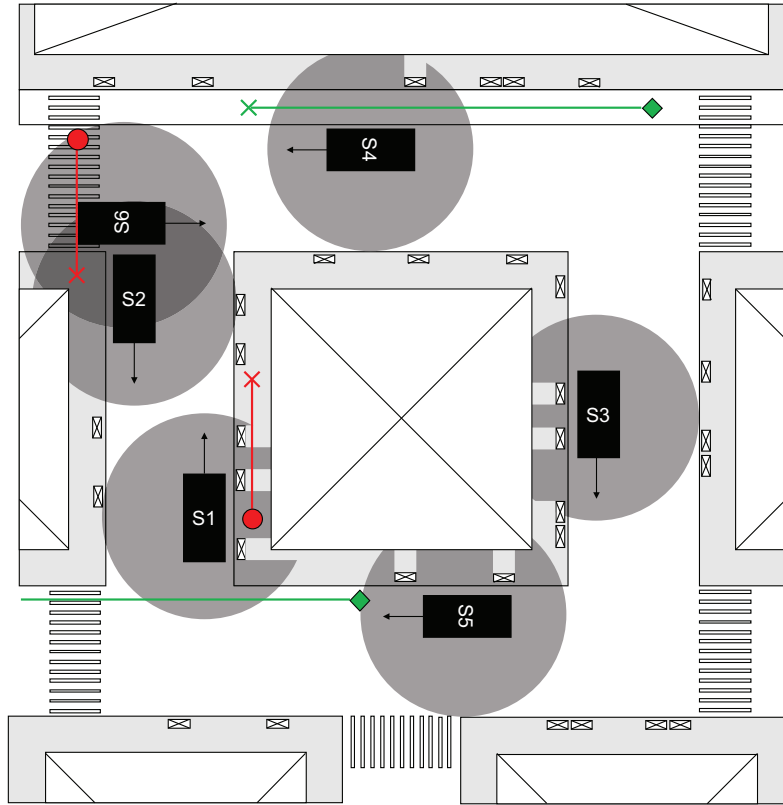
Overall, employed data collection methods are restricted to distinct areas of active mode research as each only gather certain types of data. Yet, there is no data collection method known that provides information about pedestrian and cyclist spatiotemporal properties on at different accuracy, as well as various spatial and temporal granularity, in urban settings. The need for advanced tools, and scalable systems that provide answers to spatiotemporal resolution, intrusiveness, and dynamic network conditions emerges.

## 2.3 Fundamentals

The operational objective that we address with this work is to enable researchers and policy makers to “observe active mode movements in cities at all time“. Figure 2.1 exposes the proposed active mode sensing system, AMSense, to an urban traffic setting that represents pedestrians walking on side-/crosswalk, and cyclists cycling on a designated bike path/road. In a real-world situation, however, many other traffic situations can be encountered due to the complexity of urban dynamics that continuously provoke unpredicted events. Pedestrians are usually being described with a normally distributed walking speed, whereas cyclists exhibit a greater range of speeds but more constant, and thus are being described with a wider distribution and greater mean velocity. In reality, this flexibility of speeds and directions can lead individuals to fully stop their motion, carry out less predictable movements (*e.g.*, sidestepping), or disobey traffic rules (*e.g.*, crossing at an unsignalised intersection). In addition, large intersections, multi-lane roads, and shared spaces are continuously altering the sensing system’s requirements as its mobile sensing platforms dynamically sense the environment while being in motion. Overall, following urban mobility patterns, higher densities of individuals lead to a higher sensing demand while higher densities of mobile sensing platforms lead to a higher sensing supply.

### 2.3.1 Illustrative example

To extend the technological and spatiotemporal flexibility of today’s data collection systems, we consider the following: every vehicle in a city has one or more attached sensors, and acts as a mobile sensing platform. These sensors are capable (individually or in combination) to capture pedestrians/cyclists’ mobility properties in time and space. The generated information of a vehicle’s perceived environment is filtered, and relevant data diffuses across a network of mobile sensing platforms. The information collected from each of the sensor-equipped vehicles is fused, increasing the amount of available data at a certain location, and eventually providing much greater local detailed knowledge of a city’s neighbourhood, road segment, or sidewalk, and potentially in real-time. In Figure 2.1, AMSense is exposed to an urban sensing situation in



*Figure 2.1: Urban sensing scenario in which mobile sensing platforms (black rectangles) capture pedestrians (red circles) and cyclists (green diamonds), and provide information about their presence, positions, and movements.*

which pedestrians and cyclists are observed by one or more mobile sensing platforms. The sensor equipped vehicles drive along the road network and continuously collect sensor data of their surrounding environment. This sensed data can then be processed to seek for the information of interest that is presence, locations, and movements of observed pedestrians and cyclists. These data can be shared among vehicles and communicated to a remote monitoring and control unit for additional processing, visualisation, and analyses. A comprehensive perception, in such a dynamic environment, requires the interplay of different mobile sensing platforms to obtain a detailed representation of the scene, and accurate data about sensed targets.

Collecting active mode spatiotemporal information in such a way could increase flexibility in space and time, and data could be generated at different levels of granularity and accuracy. Ideally, this sensing paradigm enables to perform different types of studies ranging from the operational level applying very detailed local data along individual pedestrians/cyclists' trajectories, over using data potentially collected with lower accuracy and aggregated to a lower degree of representation that would be applied to study macroscopic phenomena, up to studies at the strategic/tactical level using even more aggregated data over the entire network.

### 2.3.2 Requirements and performance metrics

A major trait of AMSense, is that it has to operate in increasingly unstructured environments, which are inherently uncertain and dynamic. We believe there are significant robustness and

scalability advantages when using vehicles as a mobile sensing platform that collectively achieve a global objective, while being less intrusive, and less limited by practical constraints (*e.g.*, power consumption), network reliability, and local processing capacity, than any data collection method presented earlier. For instance, vehicles are typically not affected by strict memory, processing, storage, and energy limitations, which enables the integration of various sensors, wireless transmitters and processing components. In recent years, several advanced technologies such as lightweight virtualization and edge computing have been applied to smart vehicles to enable novel applications and dynamic service deployment (Morabito et al., 2017; Laaroussi et al., 2018; Morabito et al., 2018).

We identify four prevailing elements that can generate different types of active mode data by using a network of mobile sensing platforms: network, sensing, processing, and communication. These fundamental requirements that make AMSense scalable to spatially large urban environments, and adequate for real-time implementation on rapidly moving vehicles, are illustrated as follows. Note that the idea of a system that continuously collects data on pedestrians and cyclists requires ethical, privacy, and security considerations, which shall be covered in future research. Therefore, this paper only draws the attention to such issues and advocates the need for an optional on-demand privacy component, consisting of tailored algorithms and a series of mechanisms for implementing potential privacy policies in each of the elements outlined below.

**Network** Sensor networks are typically deployed in static environments, with application specific tasks. A highly dense all-static sensor network may possibly meet the Quality of Service (QoS) (*i.e.*, service of providing spatiotemporal mobility information) requirements in an urban setting akin to Figure 2.1. However, studies have shown that sensing and networking performance of wireless sensing networks can be improved by integrating sensor mobility (Liu et al., 2005; Srinivasan & Chua, 2007). This for instance, is coupled with advantages for node deployment, and configuration strategies, supporting coverage and control in vast areas. These are required since our sensing paradigm relies upon a minimum coverage percentage (*i.e.*, percentage of area sensed by at least one vehicle). The control of a flexible temporal and spatial resolution schemes can be translated into coverage robustness as to maximise the number of vehicles sensing the same location. Furthermore, as massive data is generated by diverse sensors from each mobile sensing platform in the network, at vehicular speed, it is required to minimise the congestion probability (*i.e.*, traffic load must not surpass the bottleneck capacity of the links) when data is shared along the network. In addition, the network is required to be robust and energy efficient in routing. Meanwhile, potential applications may require data to be processed real-time, thus minimal response times (*i.e.*, latency) are required within the network. At the same time, data needs to be fresh, that is the most current status for every measurable feature, requiring low network latency from the time of sensing to the time of processing.

**Sensing** Sensing components serve as fundamental as they generate the raw data which is used later to detect, and further interpret active modes spatiotemporal properties. Different types of active mode data are required to be collected by the sensors, integrated on the mobile sensing platforms. Although different data can be gathered by diverse sensors that generate data at a variety of sampling rates and accuracy, a minimal measurement accuracy is required. Thereby, the data quality is expressed by means of data accuracy and data granularity. The latter refers to temporal and spatial resolution of the data. Robustness needs to be maintained to avoid single sensor deaths (*i.e.*, sensing robustness), and expands a vehicle's sensing capabilities (*e.g.*,

field of view). In order to generate the data, sensing elements require a minimum sensing range and power. Note that both sensing range and power have a major influence on the coverage requirements introduced earlier. The sensing application must therefore include energy-aware adaptation to dynamically adjusted sampling rates as each sensor type generates data at a different rate based on the targets or environment.

**Processing** On top of sensing, mobile sensing platforms also require embedded computational resources for data processing. A processing element is key, and can be integrated in many ways. While the processing can be performed at a centralized server in order to have a lower impact on resources, some basic filtering and anonymization of the data can also be performed before actually sending it. Especially given a proliferating amount of data will be generated, processed and stored, edge is becoming part of the processing layer besides the on-board car processing and remote cloud processing (Morabito et al., 2018). By using edge, a significant amount of processing can be offloaded before sending information to control units. Since it is infeasible to transfer the full amount of raw sensor data at all time, on-board processing and edge processing are required up to a certain extent before further transmission. To effectively perceive the dynamic movements of active modes, the sensing system requires to continuously feed real-time sensor data into the processing module (*i.e.*, frame-rate), and further understand the captured scenes (*i.e.*, processing latency) by extraction of only relevant information. Moreover, real-time processing speeds are required, as well as pre-processing procedures to reduce noise of analysed signals dynamically depending on the application. A successful implementation of such a system requires state-of-the-art algorithms.

**Communication** Data providing information on pedestrians/cyclists is required to be shareable around the network, in order to eventually be available for further analyses. Vehicular traffic is network-restricted, mainly following traffic patterns (*i.e.*, commuting hours), and travelling at maximum 60km/h in inner city-rings. Vehicles require communication capabilities that allow them to send and receive information packets to other mobile nodes, and potentially to additional control units. Some systems use short-range communication not just for detection but also for enabling collaboration between sensing platforms. As this sensing network shares time-sensitive information, data needs to be fresh (novel) and transferred in near real-time. The freshness of that data can vary, depending on how frequently a sensor is relaying readings, and is defined by sensing the correspondent application. A reliable communication is therefore required, providing answers to both potential bandwidth and latency issues. In a reliable network, the mobile sensing platforms need to transmit updates constantly, providing most current status for every sensed feature (Kadota et al., 2018). The aim is therefore to maximise the novelty of the data received at any moment, while at the same time avoiding data congestion.

## 2.4 Design of an active mode sensing network

Advances in sensing, computing and communication have driven efforts to study sensor networks, composed of a large number of densely deployed sensor nodes with self-organization, cooperative effort and on-board processing capabilities (Akyildiz et al., 2002). The vehicular based sensing paradigm we propose in this paper, is in-line with the Vehicular Sensor Networks (VSNs) for various urban monitoring tasks (Lee et al., 2006; Hull et al., 2006; Anjomshoaa et al., 2018; Minet et al., 2018). In this section, we propose the design for AMSense that cap-



tures pedestrian and cyclist spatiotemporal properties in urban areas, where a high concentration of vehicles equipped with on-board sensors can be expected, and thereby introduce the network model and its functional architecture.

## 2.4.1 Sensor network characteristics

The requirements outlined in the previous section lead to the description of main characteristics, displayed in Table 2.1, for such a mobile sensing network to be designed.

In a traditional multi-sensor system, large scale networks take in different *dynamics*: with mobile, stationary or aerial sensor nodes. We identify two distinct features for a sensing network that observes pedestrians/cyclists using vehicles as sensing platforms: i) vehicles are highly mobile, moving at different speeds along the urban street network, and ii) their mobility patterns are, to some extent, predictable due to the constraints imposed by roads, speed limits, and commuting habits (patterns). While both are typical traits of a vehicular ad-hoc network, their mobility extends the spatial and temporal scale of sensing and networking performance. Data about pedestrians and cyclists can thus be collected at many different locations within the network along the sensing platforms' trajectories, and responding to dynamic changes induced from the urban environment.

Due to mobility, the *quantity* of sensors in a network can thereby vary by several orders of magnitude over a day. As we utilize vehicles as mobile sensing platforms, and mobility patterns are strongly affected by the global mobility demand and the topology of the street network, the number of vehicles in cities follow a negative binomial distribution (Cui et al., 2017). The performance of sensing is eventually directly influenced by this characteristic, as the accuracy of estimated traffic flow variables increases with the number of operating sensing platforms observing that same scene. In addition, the spatial dimension becomes feasible, as sensing of multiple areas can occur in parallel.

Another closely tied network characteristic, with direct impact on the QoS of AMSense, is the *deployment* of sensing nodes. The sooner a pedestrian/cyclist is detected the better for the global network information quality. The optimal placement of sensors is therefore essential to meet requirements for entirely observing an area or a target. We chose to use urban vehicles (e.g., private, fleet), enabling wide coverage at reduced cost, as mobile sensing nodes to create mobile ad-hoc networks and form an inter-vehicle communication network.

*Coverage* builds upon the quantity and deployment strategy of sensors present in a network. While sensing coverage can be sparse or dense in urban settings, due to varying traffic densities, we assume nodes to be deployed with uniform density,  $q$ , subject to some temporal traffic and local patterns. The capabilities of deployed sensors in dense network vary when used individually versus when employed collectively. The coverage ratio (e.g., partial, full), with a density  $q$  that guarantees redundant coverage of an area to be monitored, depends on the sensing application and the type of data needed. For instance, besides communications and processing capabilities, network-wide microscopic sensing task (e.g., trajectories of pedestrians/cyclists in a network) can benefit from a high quantity of sensing platforms and large coverage ratio, as detailed representation of pedestrian or cyclist movements. In contrast, local macroscopic sensing (e.g., densities of pedestrians/cyclists on a road segment), requires less information, thus a limited coverage is sufficient to provide acceptable traffic estimates. Already a small

Table 2.1: Overview of designated main characteristics based on (Chong & Kumar, 2003), for a sensing network using vehicles as mobile nodes.

| Network Characteristics | Type                              | Key Performance Metrics                        |
|-------------------------|-----------------------------------|--|
| dynamics                | mobile                            | spatial resolution                             |
| quantity                | mobility demand                   | temporal resolution                            |
| deployment              | private vehicles, fleets          | deployment strategies                          |
| coverage                | spatio-temporal traffic pattern   | redundant coverage, optimal allocation         |
| composition             | heterogeneous                     | sensor fusion                                  |
| communication           | wireless                          | low latency, high bandwidth, 5G                |
| power source            | internal, on-board                | energy consumption                             |
| architecture            | multi-layer hybrid, decentralized | scalability, real-time implementation, privacy |

number of fleet vehicles could satisfy the coverage requirements (OKeeffe et al., 2019) for some specific active mode sensing tasks at macroscopic scale. We note that as traffic patterns contain distinct statistical properties, network-wide traffic state may be inferred from a learned statistical model.

The *composition* of a sensor network can be homogeneous or heterogeneous. As illustrated in the scenario of Figure 2.1, the composition of AMSense is heterogeneous as to accommodate the interplay of different sensor types.

In addition, sensor nodes need to cooperate by means of *communication* to maintain consistent real-time local information, which consists of sending and receiving data via a communication medium. This enables the sensing vehicle to exchange information with other vehicles in the mobile network. The AMSense communication medium is over wireless (e.g., 5G and Wi-Fi) to meet the low latency requirement.

The *power source* represents an influential limitation as various sensors, connectivity components, and computing equipment surely have important energy demands. For vehicle networks, a higher consumption of energy could eventually translate into reduced vehicle range. As this work presents the potential of future intelligent vehicles' "data exhaust", we will not consider issues related to continuous power delivery capabilities for the early applications.

Three different types of processing *architectures* are typically defining a mobile sensing network; centralized, distributed, or hybrid. Raw collected data needs to undergo communication and computation, in order to be further processed and aggregated to provide information about pedestrian/cyclist presence, positions, or movements. While in centralized architectures, all captured sensing data is instantly transferred to a central processing unit, this approach suffers from potential computational bottlenecks because of the sheer size of generated data, and is prone to crash in case of central unit failure or death. The inherent spatially distributed nature of multi-vehicle networks, which rely on a distributed communication and computation architecture, however, invalidate classic centric approaches. Hybrid processing architectures, in contrast, use a distributed approach to perform some level of local computation at each node, yet, still rely on a central unit to perform overall data fusion. We propose a multi-layered hybrid architecture that grounds on distributed and decentralized communication and computation, in which vehicle nodes communicate locally with surrounding vehicles. This allows the communication overhead to scale well with increase in network size and efficiently use parallel processing to process real-time data. This allows complete parallelisation of any algorithm, speed increase and proves to be a very survivable system.

## 2.4.2 Functional architecture

Figure 2.2 displays a multi-layer mobile sensing architecture with its primary functions. Remind that we make use of vehicles' sensing capabilities without interfering with their native operation. AMSense can be divided into multiple layers based on information hierarchy and computing capabilities. In a top down order, a cloud layer, on top of which diverse applications can be developed, governs an edge data processing layer that connects physical devices (i.e., mobile sensing platforms) with the cloud. Units and modules follow distinct objectives, and are represented in the figure to provide functional context.

**Cloud Layer.** At a governing level, the cloud layer (Online Fleet Operation (OFO)) offers resources for monitoring and managing the mobile sensing network, along with resources for further complex data analyses and long-term storage, while streaming information from mobile sensing platforms at any given time. This dynamic layer, manages the network, and processes less time-sensitive, non-raw data, which is already preprocessed through filtering and aggregation mechanisms by underlying layers. Information reaching the OFO can either be used for sensing operations management (*e.g.*, data acquisition) or data processing tasks. In the first place, this layer enables interactions between many possible applications, including their collaboration and data exchange, potentially in real-time. Applications are developed on top of the cloud services and provide users the possibility to further analyse the collected datasets through APIs. Those make the data available for running additional data analytics, machine learning, and visualisations on cloud servers that provide deeper insights in the collected data. In addition, the OFO offers various other services, such as coordinating user-defined sensing tasks. As such, sensing can either be performed continuously in the background or triggered by a request via cloud-based applications. Requests (*i.e.*, sensing tasks) can be defined specifically, including information about possible targets, planned routes and sensing behaviours, after which, the OFO sends task-related information to respective sensing platforms. The data collection and coordination process of those sensing tasks is performed autonomously by the cloud and underlying edge layers, based on value specific context selection that is to choose the best data sources for defined sensing tasks, without the need to contact the cloud service. A mobile sensing platform can thus wait for incoming tasks, start the data collection, and return results, or may simply publish pre-processed data without a specific request, whenever a connection is available. While the freshness of data plays a major role, data can be stored locally at the edge (*e.g.*, on-board) or globally (*e.g.*, cloud server), before or after pre-processing, and can be retrieved for further analyses.

**Edge Data Processing Layer.** Because of computational and QoS requirements and the sheer size of collected data, we move and offload computation in the proximity of mobile devices by introducing an intermediate layer responsible for additional data filtering, aggregation, processing and storage. This edge data processing layer entails the computing paradigm that delivers similar service as cloud computing, but by different means to enable a range of new benefits such as low latency, context awareness and mobility support. In addition, edge processing layer delivers similar utility computing model as cloud computing (*e.g.*, SaaS, IaaS, PaaS) but in a decentralized manner, where computing power is brought to the network's edge infrastructure. Therefore, the processing tasks on the gathered data happens at the edge, that is on mobile sensing platforms and the edge infrastructure between mobile devices and cloud services. This design reduces the load of data on the network and the cloud, as sensing, processing, aggregation and application execution are distributed over the entire network, and potentially in real-time. In contrary to a plain middle layer solution, our edge layer includes programmability and flexibility via software-defined networking (SDN) (Kreutz et al., 2015) and network function virtualization (NFV) (Han et al., 2015) to deliver ubiquitous processing capabilities across a wide range of heterogeneous hardware. For instance, the AMSense edge layer will provide image processing and raw sensing data pre-processing simultaneously. Given the heterogeneous characteristics of various instances deployed in AMSense, the edge data processing layer cooperates mutually with both cloud layer and physical layer, by acting as a bridge between elements that require dedicated interaction.

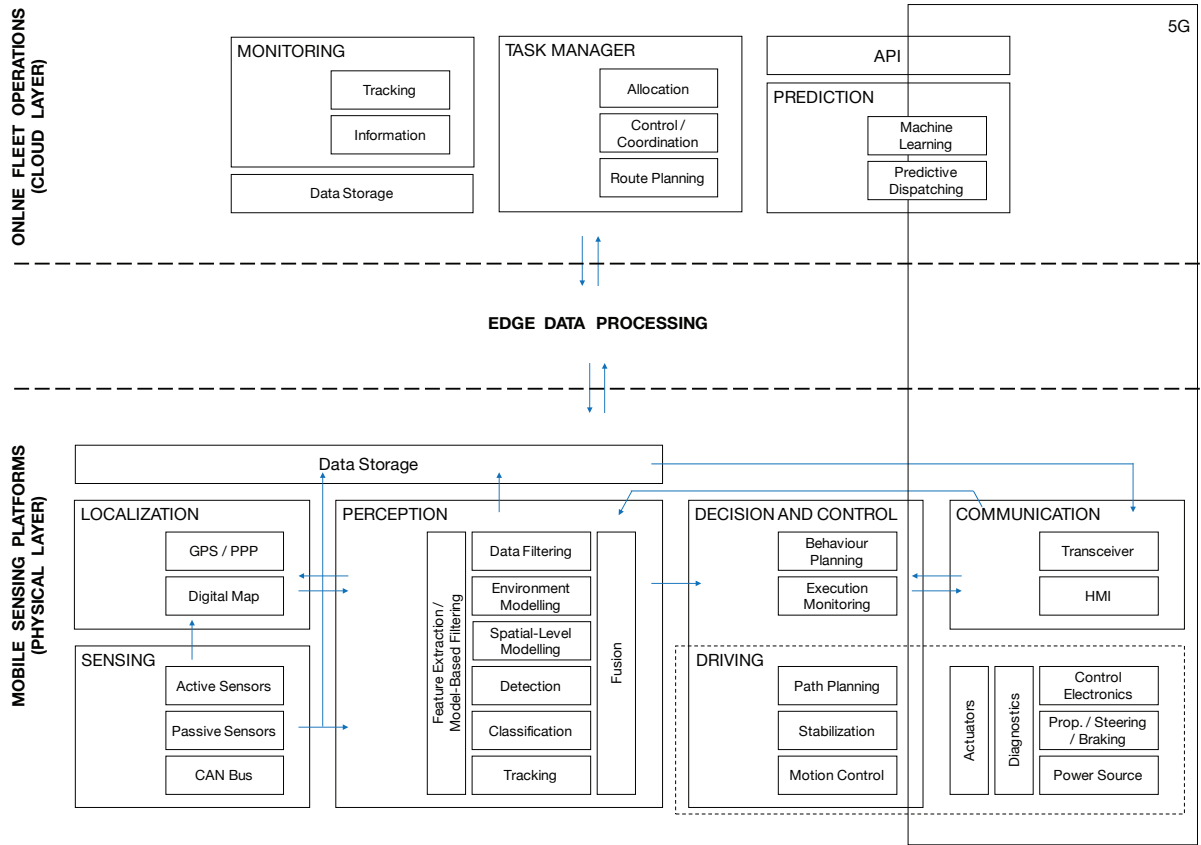


Figure 2.2: A high level, multi-layer architecture of the proposed mobile sensing system AM-Sense. Blue arrows represent data flow between different units.

**Physical Layer.** Apart from native components and functions that are necessary for a vehicle’s driving, core components of a sensing platform include sensing, localisation, processing (*i.e.*, perception, decision control) and communication.

The **sensing module** is based of sensors covering a vehicle’s external sensing capabilities, as well as internal sensors that provide information about a vehicle’s internal state. The sensing module itself remains configurable, as any vehicle integrated sensor may be activated, and thus send data, on the grounds of the different types of application. Sensor data can directly be used for localisation and map provisioning, and thus holds a common interface with the localisation module.

Although sensing and localisation are strongly linked and together form the core source of data, the latter is an independent module that has to provide two types of information on higher abstraction level. First, mobile sensing platforms are required to localise themselves (ego-localization) on a street network. The **localisation module** therefore requires to perform mapping, map updating and provides map information to other modules. Second, mobile sensing platforms are required to localise detected pedestrians and cyclists. Sensors thereby provide data input to obtain an absolute, global pose from localisation algorithms. Satellite technologies, as traditional GPS, are largely sufficient on a macroscopic level for road localisation. But shrinking down the scale to the order of a few centimetres, to perform localisation on lane- or subject-level in dynamic urban environments, localisation exceeding GPS accuracy is required. While some techniques to improve the accuracy of a traditional GPS system exists (Knoop et al.,

2012), the required stringent level of accuracy and availability for such mobile sensing application can be achieved using additional sensing sources (*i.e.*, LiDAR and camera) to produce a more accurate, robust and reliable localisation. The accuracy of an active mode's position thus depends on the technical characteristics of the sensing vehicles that capture this individual, but can also be increased with additional positional measurements that capture the individual, fused from different locations and perspectives.

The sensing module transfers raw sensor data via an interface to the ***perception module***, with additional information originating from the localisation module that follows. Different types of sensors such as LiDAR, Radar, long-range ultrasonic and forward-facing or surrounding-view camera sensors, provide depth information of a vehicle's surroundings. This data aggregate, for instance, can later be used for purposes of environment modelling. Note that sensors of each mobile sensing platform need to be considered independently, as they are subject to a vehicle's ego-motion. Vehicles have different perspectives on a scene, depending on their location and the orientation of their sensors. For these reasons, data filtering is performed to integrate one vehicle's captured scene (*e.g.*, point cloud), to another vehicle's local view, observing that same scene, at the exact same time. With detailed information about a vehicle's orientation and position, a transformation can be performed that brings the data into a universal-VSP bound coordinate system by estimating the ego motion. In other words, dynamic (*e.g.*, pedestrians) and static (*e.g.*, infrastructure) elements from a scene that are output of the mentioned sensors, can be associated to three dimensional position information, using pre-computed 3D maps of the environment. Furthermore, several algorithms (*e.g.*, feature extraction, model-based filtering) are executed to perform active mode detection, and self-monitoring of surrounding vehicles. Algorithms such as for pedestrian/cyclist detection and classification, or tracking, ground on deep neural networks. While they require substantial computing resources, their accuracy and speed is proportional to the computing resources available to them.

The ***perception module*** has interfaces to sensing, localisation, planning and control, as well as to the communication module. Perception data can directly, or indirectly, be used for broadcasting information via the communication interface to other vehicles, as to extend surrounding vehicles field of view that may be limited (*e.g.*, visibility) (Qiu et al., 2018). This perceived information is transferred to the decision and control module, where the real time map of sensed active modes and their environment is updated. Note that while the processing module is fixed, the ***storage module*** can be customised depending on application and usage.

The ***communication module*** embeds the 5G connectivity and spans across three layers, as illustrated in Figure 2.1. From architectural perspective, this implies that the communication module integrates the cloud computing, edge computing, SDN, NFV, and combines various wireless elements to deal with the requirements of AMSense services. As 5G connectivity in the future will be regarded as one of the mandatory common-pool resources (CPR) similar to water and electricity, the communication module also cater to the requirements from governmental and economical angles. This will establish a strong connection on the regulation and management in terms of interoperability, safety, cost of maintenance, public-private ownership, wireless spectrum bidding and allocation, which is necessary part of AMSense.

The ***decision and control module*** has interfaces to the perception and communication modules, as well as towards the actuators within the native ***driving module***. Algorithms in the decision and control module are the primary users of the processed information. While path planning, stabilisation and motion control are performed during the native driving of the vehi-

cle, and thus do not relate to the monitoring system, behaviour planning and execution monitoring, however, use information from perception module to potentially perform detection or tracking tasks. Those are based on messages originating from OFO interface. As mentioned before, depending on the application, sensed and perceived data can be provided at different temporal and spatial levels for active mode movement mapping updates. Behaviour planning entails operating for active mode detection or tracking where waypoints are targeted between which a route needs to be planned. Behaviour planning does however not only select the modelled movements, but also plans how it has to be executed. This manoeuvre information (*e.g.*, orientation, velocity) may be utilised by succeeding vehicles, and provided with lateral and longitudinal trajectory data to best capture the targeted active mode. For instance, the knowledge of a no-detection field (*e.g.*, occlusion), is valuable and may affect the path planning of following vehicles by changing to a lane with better view to capture a pedestrian on the sidewalk. In addition we also include execution monitoring to this module. This ensures that assigned tasks are executed as planned, and possible deviations lead to adjustments in the sensing operations. In the future, it could allow sensing vehicles to actively reposition themselves in order to optimise their sensing orientation (*i.e.*, distance, angle), using path planning that finds an optimal path when a task is assigned, while recalculating positional deviations.

## 2.5 Active mode sensing applications with urban vehicles

The complexities and edge-cases of scene perception, as well as the limitations and imperfections of sensors, make capturing the presence of active modes from a moving sensor platform full of open challenges. Using data from a single sensor source, is not necessarily sufficient to differentiate individuals from other objects in urban environments. We first categorise different sensing technologies before diving into sensor applications and examining scenario related parameters.

### 2.5.1 Sensor configuration

To start with, we classify the different sensors into those capturing the internal states of the vehicle and those capturing the states of the environment in which the vehicle operates.

The prior are described by the vehicle CAN bus, a serial broadcast bus that allow near-real-time management of most sensors and electronic devices embedded in the car. These highly integrated sensors measure steering angle, brake pressure, or acceleration rate are input for actuators related to a vehicle's native driving task. Such data may however be used to indirectly detect interactions with active modes, and hence denote their presence. In the context of driving behaviour analysis (Ly et al., 2013; Fugiglando et al., 2019), the identification of changes in a driver's behaviour could not only help recognizing hazardous situations but may also lead to describing active mode presence in real-time. For instance, think of a vehicle firmly braking in front of a pedestrian crossing (*e.g.*, zebra). The braking operation can directly be read from outgoing sensor signals via the CAN bus, while the zebra crossing may be identified based on combination of position and 3D-map. Fusing the data, we could assume that the presence of a pedestrian is likely.

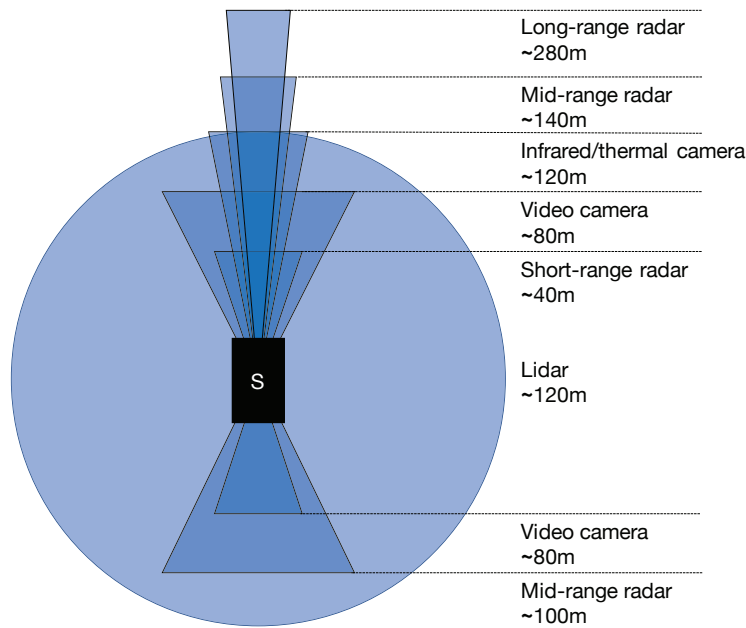


Figure 2.3: Illustration of a possible sensor configuration on a mobile sensing platform.

In a second category, we consider active and passive sensors, capturing most points in a vehicle's surrounding, as exemplified in Figure 2.3, namely: vision, LiDAR, and Radar. Vehicles equipped with vision sensors, with embedded streaming capabilities are becoming commonplace. Thereby, setups are mainly based on mono, stereo, infrared/thermal cameras. Note that mono- and stereo cameras offer a high resolution, however are subject to environmental influences, while infrared/thermal work well at all lighting/ weather conditions and raises less privacy concerns. While vision sensors can be mounted at multiple locations (*e.g.*, roof, side), the forward-facing or surrounding-view cameras are used to identify active modes and objects in the field of view, by creating a composite picture of the surrounding environment. Large improvements in computer vision based pedestrian detection were made in the recent past (Enzweiler & Gavrila, 2009; Dollar et al., 2012). LiDAR provides highly accurate long-range object detections (up to 250m), by continuously scanning the environment. They are usually mounted to the top of a vehicle, and compared with cameras can provide accurate information and larger field of view (*i.e.*, unobstructed 360-degree view) (Wang et al., 2017). While LiDAR exist at different wavelength and resolution, high-resolution LiDAR are able to pinpoint pedestrians at high accuracy and make it even possible for the system to detect human poses. Radar in contrast, output distance, velocity, and typically have high measurement accuracy. Typically, 360-horizontal coverage using short-range Radars ( $\sim 40\text{m}$  range) and mid-range-Radars ( $\sim 100\text{m}$  range) can be achieved with 4-6 Radars mounted at each corner/side of the vehicle. As Radar is less affected by external conditions, they provide redundancy for camera and LiDAR.

The detection accuracy and correct localisation can be significantly improved by fusing data from a set of sensors. This for instance, supports the use of a sensing network with multiple sensors, while making use of vehicles' movements to increase detection and measurement performance through the extension of different field of views. This capability that provides depth perception of the surrounding environment is likely to become pervasive in future vehicles. Note that the mentioned sensor components all depend on different system integration and amount of processing.



## 2.5.2 Active mode sensing scenarios

We further give an overview of urban traffic situations, illustrated in Figure 2.4, where pedestrians and cyclists can be sensed with AMSense. Four examples illustrate interactions between mobile sensing platforms and active mode behaviours in typical urban settings: a) target is in field of view, b) is not in field of view, c) unclear situations due to occlusion, and d) the target deviates from his linear trajectory (*e.g.*, crossing). Note that we assume active modes not to be restricted by the pedestrian or cyclist network (*e.g.*, shared spaces), while vehicles to only use the underlying street network.

A pedestrian or cyclist may be identified or tracked in time and space by one or more sensing platforms. However, when assigned an ID (*e.g.*, “pedestrian 1”) and tracked for a certain time interval, the tracked target may lose its temporary ID in situations where no target-sensing is possible (*e.g.*, occlusion, location not sensed). The target will most likely be assigned a new ID, when re-entering the field of view of that same (or different) vehicle. While a vehicle sensing a target can solely be expressed as the estimate detecting that target, the quality of collected data depends upon the interplay of the sensing system’s technological capabilities and actual urban traffic conditions. Overall, the granularity and accuracy of the data are typically determined by sensor types, while the accuracy also depends upon sensing positioning, node coverage and external conditions (*e.g.*, weather).

In the first scenario (a), the pedestrian and cyclist are in unobstructed field of view to their closest vehicle, and can thus directly be captured by that vehicle. This situation occurs when the target is in detection radius of the sensing vehicle, and in the case that no additional active mode, or object hinder a clear view on the target. In this situation we assume that data about that captured individual will always be available by means of at least one sensing vehicle, whether it is in motion or not. Collected information about that pedestrian/cyclist could therefore include presence, locations, speeds, and movements, generated at each sampling time until moving out the detection radius. The measurement accuracy of that individual, is affected by the sensing vehicle’s position, that is the distance and angle to the target. Furthermore, the effects of such sensing network on data quality becomes apparent when an active mode is captured by more than one sensing platform, as in Figure 2.4 vehicles 2 and 6. Both generated data are fused (data needs to be fresh), extending the vision on that pedestrian while potentially increasing the accuracy and granularity of the data.

In a second scenario (b), the pedestrian and cyclist are in unobstructed field of view to their closest vehicle, however can’t be directly captured as they are not inside the vehicle’s detection radius. This situation occurs in case of large distances or wide angles between the target and sensing vehicle. In this situation no data on the present active modes is generated. Although vehicles may sense their surroundings, no pedestrian/cyclist is detected until they reach one of the vehicles’ sensing radius. Yet, prior data that was generated at the time a pedestrian/cyclist was moving in a vehicle’s sensing radius, might still be available in the network (*e.g.*, data storage). Almost fresh information (*e.g.*, near real-time) may then be used to estimate an active mode’s position for a limited time interval. This situation can be altered by increasing the coverage of that sensed area.

In a third scenario (c), the pedestrian and cyclist are in obstructed field of view to their closest vehicle, and thus can’t be sensed by that vehicle, in the time of sampling. This situation occurs when an object (*i.e.*, infrastructure, nature) hinders the view, and thus makes perception

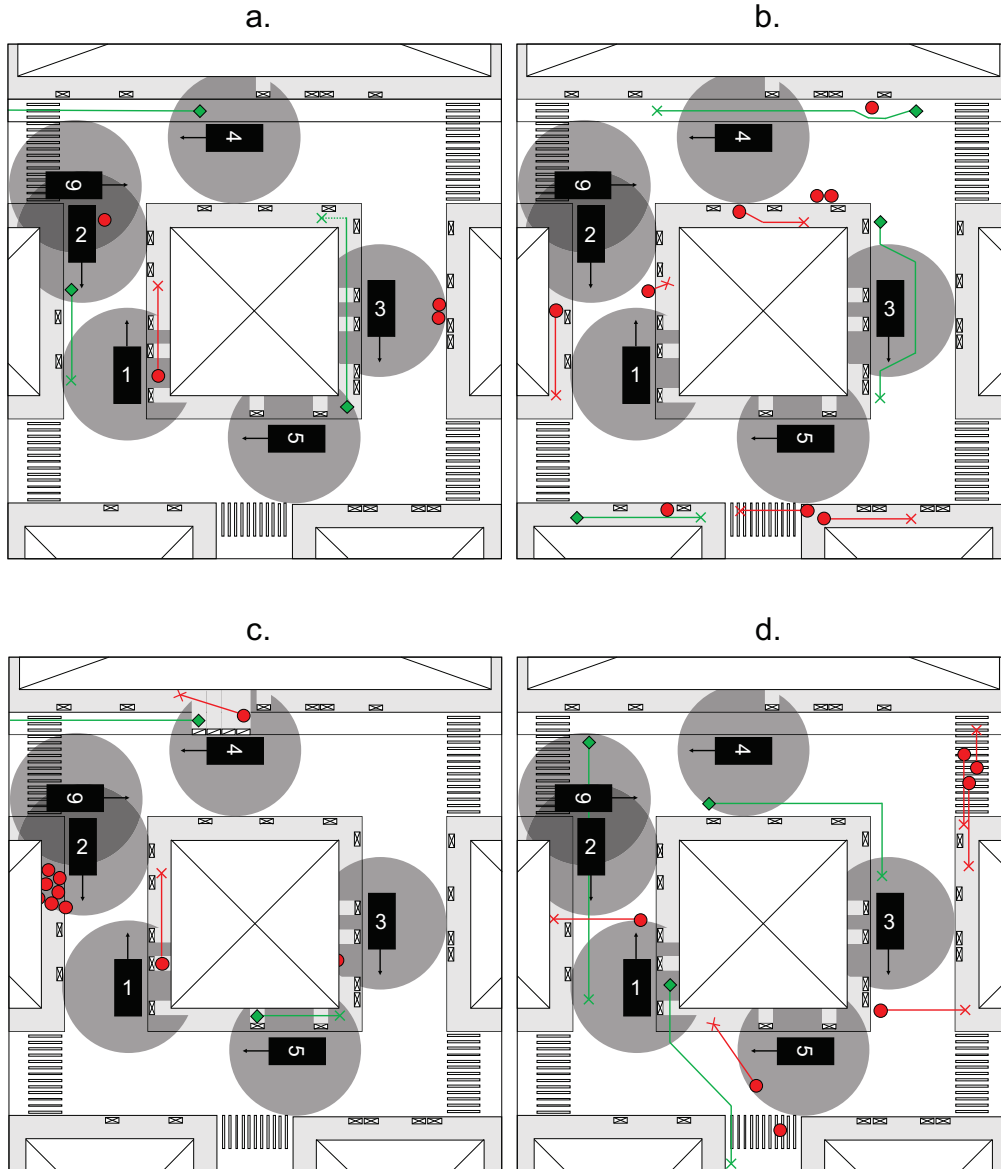


Figure 2.4: Overview of potential traffic situations, where mobile sensing platforms sense their environment while pedestrians (red circles) and cyclists (green diamonds) are a) in field of view, b) not in the field of view, c) not or partially detectable, d) crossing.

of the active mode impossible, even though the target is in radius of detection. This situation also occurs in settings where crowds make individual detection more difficult. Full occlusion entirely hinders collecting any data about that individual's presence, location, speed, or movements, until reappearing in the same (or another) vehicle's field of view. Partial occlusion, yet, reduces dramatically the quality of data, however, data about that active mode may be processed up to a certain extent. The positioning of a mobile sensing platform and its sensors influences the occlusion shadow, and therefore, the amount of collected information. Thereby, data fusion enhances the overall detection and perception, as diverse sensors could capture targets differently. At the same time, increasing the sensing coverage, is likely to expand the collective field of view on a scene.

In a fourth scenario (d), we show a subset of different crossing situations in urban traffic settings. Active modes may use signalised or unsignalised locations to cross a street. Involved vehicles adapt their ego-motion to the crossing situation, which shall have no effects on sensing capabilities. However, less vehicles might be able to see that crossing individual as they'll have a different perspective on the environment (parallax), *e.g.*, crossing in front of a vehicle create an occlusion shadow for the following vehicles. Remember that perception data from the first vehicle can then be used to extend the field of view of following vehicles.

The quality of collected data is, technological capabilities and external conditions aside, mainly influenced by the positioning of sensing vehicles to their targets, as well as the sensing coverage at a certain location. Likewise, the movement of pedestrians and cyclists is equally influencing the collection of data quality, due to occlusion introduced by infrastructure, vehicles or groups of people. In addition, the movements of vehicles in relation to a sensed radius influences the granularity of the data. Remind that, collecting datasets about labeled individuals across the network, as for instance needed in route choice studies, seems impossible without using additional sensing technologies (*e.g.*, Bluetooth).

### 2.5.3 External operating conditions

Although we cannot account for all environmental factors that might affect the sensing performance, we describe the ones that are most likely to adversely affect the correctness of sensors, protocols, and algorithms. In the remainder of this section, we describe weather effects, occlusion effects, and noise model.

In addition to dealing with internal system failures resources (*e.g.*, bandwidth), the mobile sensing network, and more particularly single sensor nodes may fail or be blocked due to physical damage, or environmental interferences. Environmental reliability, denoted as  $R_k(t)$ , can be modelled using the Poisson distribution and represents the probability functioning without interruption during sampling time  $t$ . Furthermore, weather conditions (*e.g.*, low visibility), can directly affect the active mode detection, and cause false positives/negatives. In scenarios with low lighting conditions, we expect a drastic decrease in sensing performance of most vision sensors (infrared excluded). Temperature is not expected to have a direct effect on overall performance.

Road conditions (*e.g.*, street quality) and the presence of obstructing infrastructure or urban greenery can have a negative impact on the quality of sensing targets and communication between nodes. One of the main concerns when detecting pedestrians and cyclists is occlusion, which occurs especially in high density scenes or when objects hinder a clear view in an observed area. A vehicle's field of view towards the sidewalk may be occluded by trees, parked cars or urban infrastructure. This prevents the mobile sensing platform to correctly observe the actual scene, and thus missing out possible pedestrians (or cyclists). Although some areas might present high risk of occlusion, this occlusion might not be present for very long since the active modes and sensing vehicles are moving at different relative speed. In an eventual future, in which one could think of vehicle manufacturers or fleet operators allowing the placement of sensors at various locations on a vehicle, the surrounding environment could be perceived from a multitude of distinct points on a mobile sensing platform. It should be noted that vehicle-free areas make up a sizeable fraction of most urban areas. Active modes would therefore have to be

monitored with other mobile vehicles (*e.g.*, bikes, drones), or static sensors that are equivalent to mobile sensing platforms at speed 0, and their fusion would be worth considering in future studies.

Typically, such a mobile sensing network is subject of noise-corruption that increases with the quantity of sensor nodes involved. As mentioned earlier, the detection of a pedestrian/cyclist is only expressible as a probability. All raw sensor data come with uncertainties (*i.e.*, sensor efficiency, weather, road conditions) and hence inherit noise. Probabilistic methods, such as Kalman filter, addresses this margin of error.

## 2.6 Conclusion

In this paper, we presented AMSense, a novel mobile sensing system that uses connected multi-sensor equipped vehicles to build a sensing network which captures pedestrians and cyclists spatiotemporal properties in urban areas. The collected data about pedestrian/cyclist presence, locations, and movements can be used as input for a variety of studies that require active mode information at diverse macro- and microscopic levels. Future work will dig into challenges of reliability, scalability, as well as ethical, privacy, and security considerations. Moreover, future work will investigate dynamic and context-aware data collection, explore the potential of measuring spatiotemporal densities, speeds and flows using such a mobile sensing system, and study active repositioning of sensing vehicles to optimise vehicle allocation.

## Chapter 3

# Network-constrained tracking of cyclists and pedestrians

---

Chapter 2 has identified data and technical requirements distinct to active mode research and practice within the new sensing paradigm. In the current chapter, we zoom in on pedestrian and cyclist observations from types of IAS, especially on how to make sense of these observations in time and space.

We do so by assuming the availability of data with a minimal set of features as to show how little needs to be sensed to obtain a comprehensive microscopic and macroscopic picture of traffic flows in a larger area. This chapter proposes a new method for advanced traffic applications, tracking an unknown and varying number of moving targets (*e.g.*, pedestrians or cyclists) constrained by a road network, using mobile (*e.g.*, vehicles) spatially distributed sensor platforms. We showcase the framework with a simulation study, highlighting how network-constraints inherent to transportation infrastructure is beneficial when estimating the number and location of pedestrians and cyclists. The proposed framework can conceptually be extended to any type of sensing realm and presents a versatile tool for different stakeholders in various contexts. This chapter also directs attention to latent privacy concerns for potential applications.

This chapter is published as a journal article: Vial A., Hendeby G., Daamen W., van Arem B., and Hoogendoorn S.P. (2022). Framework for Network-Constrained Tracking of Cyclists and Pedestrians. *IEEE Transactions on Intelligent Transportation Systems*.

---

### 3.1 Introduction

Inferring the number of pedestrians and cyclists, as well as their individual states (*e.g.*, position, velocity) from a sequence of measurements, allows drawing a complete micro- and macroscopic picture of the traffic flows in an observed area. This is of great value for advanced surveillance applications in traffic operation, control and management.

Traditional setups for these applications rely on stationary or participatory sampling technologies. Stationary sensing systems gather data that can be used to reconstruct trajectories or counts, but only at local scale (*e.g.*, cross- or very short road sections). At the same time, crowd-sourced data from mobile and other wearable devices allow tracking individuals through a network, but require direct or indirect collaboration of the tracked individual and are mostly sparse in nature (not all persons are detected). Current advanced surveillance applications and traffic monitoring typically consider longer time horizons to associate new measurements with existing tracks to monitor individuals over time. Proposed methods mainly consider stationary sensor settings, which often produce a visual stream from a fixed location (Barthélemy et al., 2019; Duives et al., 2020), however come with low spatial resolution and scalability.

With novel distributed-computing and connectivity capabilities, a growing number of on-board sensors, the collection of data on pedestrians and cyclists is expected to further intensify with the advent of connected autonomous mobile systems. Self-driving vehicles, drones, or other types of connected robots will enter populated environments and may act as mobile sensing platforms generating a proliferating amount of data about the platform's internal state, but also about the static and dynamic local area they observe. With the collective intelligence and innate mobility of such sensor platforms, pedestrians and cyclists traffic characteristics could be captured at an extended spatial and temporal scale (Vial et al., 2020). In the future, knowledge about position, motion state, and pose of people could enable next generation traffic or crowd surveillance systems to estimate the number of people and reconstruct trajectories across the network.

Yet deriving complete trajectories of indistinguishable pedestrians and cyclists in a network using observations from mobile spatially distributed sensor platforms, is complicated. First, there is an unknown and varying number of people, where position and motion states of individuals are unknown. At the same time, noisy sensors, changing environmental conditions, or occlusion are responsible for missed detections and clutter. And with no *a priori* information about which observations originate from which existing or newly detected individual, the many possibilities of assigning a measurement to an individual complicate the task.

There have been major developments in pedestrian and cyclist detection and tracking, for instance, in areas related to autonomous navigation and control. Safety-critical applications, *e.g.*, self-driving vehicles rely on accurate human motion prediction and path planning (Rudenko et al., 2020; Keller & Gavrilu, 2014; Kooij et al., 2019). Proposed methods and pilot deployments however typically consider spatially restricted short time horizons. Thereby presence and state information is crucial to better understand and anticipate actions, that is knowing what an individual will do next, *e.g.*, start, continue, or stop walking.




Figure 3.1: Illustration of the concept showing a simple road network and the underlying Markov Process.

In the last decades, different target tracking approaches have been proposed that aim at estimating the number and states of (multiple) dynamic objects using noisy sensor measurements, taking source from the tracking community. Existing algorithms, however, consider short time horizons and are mostly designed for unconstrained motion in two- and three-dimensional space, where complexity tends to grow exponentially due to the number of hypotheses. Attempts have been made to better model the behaviour of pedestrians, *e.g.*, (Batkovic et al., 2018), which aims at learning to predict for collision avoidance, and (Luber et al., 2011), in which spatial motion model is learned from observations. Due to their added complexity, these methods fit poorly into standard multi-target tracking methods.

Road information has previously been used to support ground target tracking. The approach proposed in (Pannetier et al., 2004) tracks the objects in free space and then, in a wide sense, projects the estimates onto the road network. This ensures that targets are estimated on the road, but does not improve projections as done in this paper. In (Kirubarajan et al., 2000), a *variable structure interacting multiple model* (VS-IMM) is introduced, which uses different motion models based on the target location to keep the objects on the roads. In (Ulmke & Koch, 2006), targets are described in a combination of global 2-D coordinates, and the quasi 1-D coordinates as used in this paper. This allows for more efficient predictions while, as in (Kirubarajan et al., 2000) and (Song et al., 2018), the VS-IMM is used. Compared to the proposed method, the number of target hypotheses are regularly reduced, keeping them alive until observations can determine which one is correct in the proposed *multi-hypotheses tracker* (MHT). Furthermore, (Kirubarajan et al., 2000) and (Ulmke & Koch, 2006) consider airborne ground moving indicators as sensor, which is different from ground based sensors moving on the road network as introduced in this work. In more recent efforts, (Song et al., 2018) uses measurements from static sensors and include road map information, however, focus on tracking interacting vehicles rather than on the effects of road constraints when tracking targets in a network. (Zheng & Gao, 2018) extends the idea of including road map knowledge to random finite set (RFS) methods, however, using a more involved road representation and handling than the method proposed in this paper. (López-Araquistain et al., 2019) also uses a different road and coordinate representation that requires projections, however, the concept of branching hypotheses in ambiguous situations (*e.g.*, at intersections), also found in (Krishanth et al., 2014),

is similar to the handling found in our approach.

The key contribution in this paper is to introduce the concept of network bound targets into the multi-target tracking problem. This is done by introducing a target representation, comprising a traditional target tracking representation and a discrete component placing the target on a given edge in the network. This model is then used to derive a *network-constrained multi-hypotheses tracker* (NC-MHT). The inclusion of knowledge about the network structure allows for more efficient target predictions over extended periods of time and simplifies the measurement association process, as compared to not utilizing a network structure. The network bound target model is derived with tracking of an unknown and varying number of pedestrians and cyclists in an urban setting using mobile spatially distributed sensor platforms in mind. Hence, the NC-MHT is evaluated on three simulations of targets moving around on an urban infrastructure of connected roads, highlighting different properties introduced by adding the network constraint. The NC-MHT opens up the field for new applications for network-wide traffic surveillance, using information about the number of individuals and their states, to enhance advanced traffic operation, control and management systems.

The remainder of this paper is organized as follows. The next section presents relevant background theory on fundamentals, namely Bayesian filtering and multiple target tracking. Section III introduces the mathematical problem formulation and network-constrained system models. Section IV presents the derived NC-MHT filter and details the integration of the network structure. Section V presents simulation results from three tracking scenarios. Section VI discusses the outlook and potential extensions of the proposed approach. Conclusions are drawn in Section VII.

## 3.2 Background

Using conventional methods for multiple target tracking (MTT), the individual targets are tracked using Bayesian filters (Fang et al., 2018), following a step where the available observations in each scan are assigned to the different tracks. In this context, a track represents a potential target and contains information about its estimated past and current state, and a scan is a set of observations received at the specific point in time. A MHT considers several different association hypotheses in parallel. Additional logics handle track creation and termination of tracks over time. The steps are outlined here, and further adapted to network bound targets in the following section.

### 3.2.1 Bayesian filtering

The state  $x_k$  (e.g., position and velocity) of a target at time  $k$  can, given the observations from time 1 to  $k$ ,  $z_{1:k}$ , be estimated using recursive Bayesian filtering. This is achieved using a two step iterative process that sequentially predicts and updates the probability density function of the target state  $p(x_k|z_{1:k})$ ,

$$p(x_{k+1}|z_{1:k}) = \int p(x_{k+1}|x_k)p(x_k|z_{1:k}) dx_k \quad (3.1a)$$

$$p(x_k|z_{1:k}) = \frac{p(z_k|x_k)p(x_k|z_{1:k-1})}{p(z_k|z_{1:k-1})}. \quad (3.1b)$$



The first equation predicts future states, given the current state and a motion model  $p(x_{k+1}|x_k)$ , and the second step incorporates information from a new observation  $p(z_k|x_k)$  into the estimate.

For linear models with Gaussian noise, where both the dynamic and measurement model are linear, the seminal *Kalman filter* (KF) (Kalman, 1960) provides the analytic solution to the Bayesian filtering recursion. For the case that nonlinear models are used, the *extended Kalman filter* (EKF) (Jazwinski, 1970), or *unscented Kalman filter* (UKF) (Julier & Uhlmann, 1997) can be used to approximate the solution, and in highly nonlinear scenarios the *particle filter* (PF) (Ristic et al., 2004) can be used.

### 3.2.2 Multiple target tracking

Multiple target tracking is an extension to the state estimation problem where both the number of targets and their states should be estimated, based on available scans. The observations that make up the scans are noisy, a target may sometimes fail to produce a detection, and clutter (false observations) which all complicate the task.

The core of the MTT problem lies in the association of observations to the right track, as given the association, the state of the tracks can be estimated using Bayesian filtering. Data association is needed as it is not known which observation originates from which target, not to mention which observations are clutter. Classical MTT solutions as the *global nearest neighbour* (GNN) tracker, the *joint probabilistic data association* (JPDA) filter, and the *multiple hypothesis tracker* (MHT) (Reid, 1979) differ in the way the association is performed and how many hypotheses are maintained. See (Blackman & Popoli, 1999; Bar-Shalom & Li, 1993) for more details. More recently developed MTT methods, *e.g.*, the *probabilistic hypothesis density* (PHD) filter (Vo & Ma, 2006; Krishanth et al., 2017), the *labelled multi-Bernoulli* (LMB) filter (Reuter et al., 2014), and *Poisson multi-Bernoulli mixture* (PMBM) filter (García-Fernández et al., 2018), are based on random finite set statistics and are derived differently, resulting in slightly different methods. None of these methods, however, are designed to utilize network constraints.

The MHT is one of the most popular and common MTT methods. The use of several association hypotheses in parallel took off with Reid's seminal paper (Reid, 1979), yet requires many approximations to become a computationally tractable technique as the number of possible track associations increases exponentially with time. Since then, it has been further developed as this technique gained some momentum in the tracking community, and has been applied in specific application domains. The MHT technique relies on evaluating the probabilities of sequences of measurements from various targets. Two main variants of MHT can be found in literature: Hypothesis-Oriented (HO-) (Reid, 1979), and Track-Oriented MHT (TO-MHT) (Kurien, 1990). Cox (Cox & Hingorani, 1996) further contributed to the design of an efficient MHT implementation by introducing Murty's algorithm (Murty, 1968). The latter reduces the computational complexity as it helps avoiding considering unnecessary hypotheses.

### 3.3 Problem formulation

This paper considers a regular MTT problem, with an additional network constraint as imposed, *e.g.*, by targets being bound to a system of roads. That is, given a set of mobile sensors, with limited field of view, determine the number of targets present in the tracking volume and their respective state. Standard target tracking assumptions are assumed to apply:

- A1: Targets act independently, without influencing one another.
- A2: The number of new targets at each time is Poisson distributed, with the rate  $\lambda_{NT}$ , and targets appear uniformly in the tracking volume.
- A3: In a scan of observations from a sensor, each observation either originates from a single target or is a false detection, and each target produces at most one observation.
- A4: The number of false observations (observations not originating from a target) in a scan is Poisson distributed with the rate  $\lambda_{FA}$ , and the observations are uniformly distributed in the sensors field of view.
- A5: The probability to detect a target that is within the field of a sensor is constant,  $P_D$ .
- A6: The probability a target survives from a time to the next is constant,  $P_S$ .

Additionally, *this paper assumes targets to be network constrained*, that is:

- A7: The targets are bound to network constraints, and hence cannot move freely in the tracking volume. A target is always associated with exactly one discrete network state at any time.

Here, it will be assumed that the network constraint is in fact a road network. However, the theory developed applies to any situation where the targets are associated to a discrete state that can be modelled as a Markov process. Notation will be kept as general as possible.

#### 3.3.1 The network constraint: road network

Targets are constrained to move on roads, sidewalks, and cycle paths, which make up the considered network constraints defined in Assumption A7. A key observation is that a target can only be on one road segment at the time, and once the target reaches an intersection it continues on to another connected road segment, which limits the motion of the targets. This can be modelled as a discrete Markov chain, where the road segments make up the states and the transition probabilities<sup>1</sup> model the probabilities to transition from one road to another once the target reaches an intersection. This can be represented with a directed graph  $G = (V, A)$ , with a set of vertices  $V$  (representing the different states or road segments) of size  $n$  and edges  $A \subseteq V \times V$  (the possible transition from one road segment to another) of size  $m$ . A weight function  $p : A \rightarrow [0, 1]$  assigns a non-negative probability  $p(\delta'|\delta)$  to transition from state  $\delta$  to segment  $\delta'$  once the target reaches the end of the road segment represented by  $\delta$ . It follows from the rule of total probability that  $\sum_{\delta'} p(\delta'|\delta) = 1$ . Furthermore, each vertex contains information about the length of the

<sup>1</sup>A discussion on the determination of the transition probabilities is provided in Section 3.6.2.

road segment it represents,  $L(\delta) > 0$ . The graph hence represents a topological map of the road network.

### 3.3.2 Network-constrained state-space model

To utilize the road constraints, a discrete component is added to the target state to indicate which road segment the target is on, and the “regular” part of the state will be used to track the target within the node. Both the prediction and observation model are modified to take the discrete component in consideration.

Targets evolve independently of each other, in discrete time, and the state of each target is represented by a vector

$$X_k = \begin{pmatrix} x_k \\ \delta_k \end{pmatrix},$$

where  $x_k$  is a real-valued state vector describing the target motion within a node, and  $\delta_k$  is the state in the Markov process. The semantics of the state space  $x_k \in \mathbb{R}^{n_x}$  can vary depending on the considered application, yet usually includes components related to the target’s kinematic state such as position, velocity, and possibly acceleration, along the road segment, but could also be used to describe the lateral position on the road. In the end, the actual position of the target in the world is a combination of the current node  $\delta_k$ , giving the general position determined from map, and the  $x_k$  component which specifies the position along the road segment as illustrated in Fig. 3.1.

#### Motion model

The nonlinear dynamic model describes the motion within a node as

$$x_{k+1} = f(x_k) + v_k, \quad v_k \sim \mathcal{N}(0, Q_k) \quad (3.2)$$

where  $v_k$  is process noise that is assumed to be white and Gaussian with covariance matrix  $\text{cov } v_t = Q_k$ . The discrete state  $\delta_k$  follows the transition probability  $p(\delta_{k+1} | \delta_k, x_{k+1})$  where  $\delta_{k+1} = \delta_k$  unless  $x_{k+1}$  indicates that the target reached the end of the current road segment. Once a state transition occurs, the node specific state, and distance travelled along the current road segment,  $x_{k+1}$  is updated accordingly. Typically, this means resetting the distance travelled. In the considered application, when a transition takes place and  $x_{k+1}^p$  is the position component of the continuous state, remove the distance travelled on the previous segment to reach the end of it  $x_{k+1}^p \leftarrow x_{k+1}^p - L(\delta_k)$ . It could, however, potentially also include additional adjustments as a result of specific properties of the new road segment. The adjustment procedure must be repeated if after compensation the target is past the end of the next segment too.

#### Observation model

At any time instant  $k$ , sensor platforms may deliver scans,  $Z_k$ . The position of the mobile sensor is assumed to be known, and as defined in Assumption A5, each target is detected with

probability  $P_D$  when in the field of view of a sensor. The state variables above are related to measurements  $Z_k$  according to a nonlinear observation model

$$Z_k = (z_k, z_k^\delta(\cdot)) \quad (3.3a)$$

$$z_k = h(X_k) + e_k, \quad e_k \sim \mathcal{N}(0, R_k) \quad (3.3b)$$

$$z_k^\delta(\delta') = \Pr(\delta_k = \delta') \quad (3.3c)$$

where  $h(x_k)$  represents the state of the target and relates to an ideal (vector valued) measurement, and  $e_k$  is the measurement error which is assumed to be Gaussian with covariance matrix  $\text{cov}(e_k) = R_k$ . Furthermore, each observation includes the probability that the target was observed in the discrete state  $\delta'$ ,  $z_k^\delta(\delta')$  — a simplification, reasonable in many cases, is that the discrete state information is accurately provided, *i.e.*, that only one state has non-zero probability.

Because several sensors can operate simultaneously, let  $S = \{s_v\}_{v=1:V}$  be the set of mobile sensors, where  $V$  is the number of operating sensors that acquire data of dynamic targets present on the road network. A scan  $Z_k^v$  is the set of all  $M_k^v$  measurements received by a sensor platform  $v$  at time  $k$  such that  $Z_k^v = \{Z_k^{v,1}, Z_k^{v,2}, \dots, Z_k^{v,M_k^v}\}$ , and where  $Z_k^{v,i}$  represents the  $i$ th observation received at scan  $k$ , by sensor  $v$ . This way, a target may be observed by multiple sensor platforms at the same time instant.

### 3.3.3 False observations and new targets

As part of the point-to-object and defined in Assumption A3, each target gives rise to at most one measurement per sensor, at each time step. Targets in the field of view are assumed detected and observed with probability  $P_D$  (Assumption A5). Further, the number of new targets  $n_{NT}$  and false observations  $n_{FA}$  in the volume  $V$  are assumed Poisson distributed,

$$p(n, \lambda, V) = e^{-\lambda V} \frac{(\lambda V)^n}{n!}, \quad (3.4)$$

with intensity  $\lambda_{NT}$  and  $\lambda_{FA}$ , respectively, see Assumptions A2 and A4. Both new targets and false observations, together denoted extraneous observations, are assumed uniformly distributed in the current tracking volume. The total spatial density of the extraneous observations also follows a Poisson distribution with intensity  $\lambda_{EX}$  and can be expressed as follows

$$\lambda_{EX} \triangleq \lambda_{NT} + \lambda_{FA}, \quad (3.5)$$

which, as will be seen, can be used to simplify some expressions.

**Note:** In the network constraint cases, the sensors can be assumed to only produce state (road) bound observations, which is beneficial compared to the general case where no such restrictions exist.

## 3.4 Network-constrained multiple hypothesis tracking

In a MHT, several potential targets are tracked using separate single target tracking filters, and a higher level logic decides which observations to pair with the currently maintained tracks,

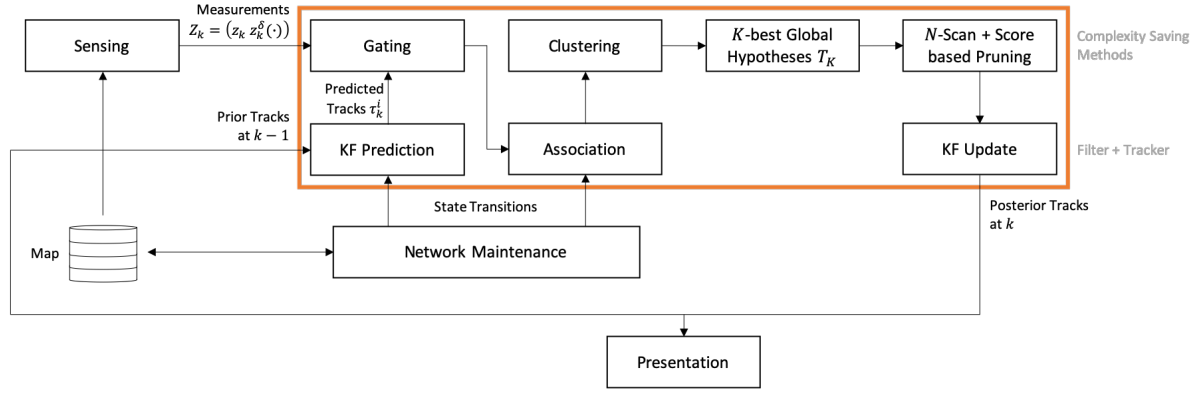


Figure 3.2: Outline of proposed tracking framework, with traditional MHT logic inside orange area.

which observations should be considered to be false, and when to create new tracks. In the MHT, several such hypotheses are maintained in parallel, while at the same time estimating the probability of each of the different hypotheses to be correct. Key components are to generate the appropriate hypotheses, and to reduce the number of considered hypotheses as much as possible to keep the computational complexity as low as possible. The network constraint, being discrete in its nature, simplifies the hypothesis handling. This section outlines the MHT formulas adopted to take the network constraints into consideration. Fig. 3.2 provides an overview of the different elements of the proposed algorithm.

### 3.4.1 Extended Kalman filter

The EKF (Jazwinski, 1970) is arguably the most common technique used for single target tracking in the MHT. The EKF does not deal with discrete components, hence it cannot be directly applied to the hybrid state proposed in Sec. 3.3.2.

#### Time update

The solution is to condition on the discrete state  $\delta_k$ , and use the regular EKF for the continuous part of the state  $x_k$ . This yields, for a given  $\delta_k$ , for the time update

$$\hat{x}_{k|k-1} = f(\hat{x}_{k-1|k-1}), \quad (3.6a)$$

$$P_{k|k-1} = F_k P_{k-1|k-1} F_k^T + Q_k. \quad (3.6b)$$

where  $\hat{x}_{k|k-1}$  is the mean of the estimate of the state at time  $k$  given the measurements up until time  $k-1$ ,  $P_{k|k-1}$  is the matching error covariance matrix, and  $F_k = f'(\hat{x}_{k-1|k-1})$  the linearization around  $\hat{x}_{k-1|k-1}$  of the dynamic model.

Caveat, performing this prediction step might position the target past the end of the current discrete state. The interpretation in the described setting is that the target reaches the end of the current road segment. That is, the target should transition to a new discrete state. This is done following the underlying Markov model as described in Section 3.3.1. If there are several different possible transitions, where each have transition probabilities, all are taken and new

track hypotheses are created, and the hypothesis probability is updated accordingly by the MHT logics, as described below. After the transition the continuous state  $x$  is compensated in each transition, as described in Sec. 3.3.2. As for the case there is no discrete state to be propagated to, *e.g.*, a target is existing the tracking volume of interest, the target is simply removed from the set.

### Measurement update

When an observation is obtained, again conditioned on the discrete state, the continuous part of the state can be updated, provided the discrete state is possible in the observation, using

$$\hat{x}_{k|k} = \hat{X}_{k|k-1} + K_k(z_k - \hat{z}_k), \quad (3.7a)$$

$$P_{k|k} = P_{k|k-1} - K_k H_k P_{k|k-1}, \quad (3.7b)$$

$$\hat{Z}_k = h(\hat{x}_{k|k-1}) \quad (3.7c)$$

$$S_k = H_k P_{k|k-1} H_k^T + R_k, \quad (3.7d)$$

$$K_k = P_{k|k-1} H_k^T S_k^{-1}, \quad (3.7e)$$

where  $H_k = h'(\hat{x}_{k|k-1})$ . If no observation is available, the update step is simply skipped, and if several are available the step is repeated.

Contrary to the prediction step, the observation update will never trigger transition in the discrete state, if the simplified version is used. The update is performed under the presumption that the observation places the target in the current discrete state, hence the estimate stays in the same state. The only ambiguities in this step is due to different observations being associated with the current track.

### Corner cases

The above approach assumes that each track hypothesis can fully be assigned to a discrete state. This is often a reasonable approximation as long as the estimated track is far from the end of the current road segment, as compared to the uncertainty of the estimate. However, if the track is close to the edge of the road segment or the uncertainty is considerable, with a non-negligible probability the track could instead be on an adjacent road segment. Two different solutions to this are considered here: always splitting tracks with significant overlap, and splitting when needed for proper association with observations.

In the former case, when a track is predicted close to an intersection, split the track in  $n - 1$ , where  $n$  is the number of potential roads at this intersection, and divide the probability according to the probability mass on all segments other than the segment the target is currently on (based on the uncertainty). The resulting two new tracks should have means combining to the original mean, and a combined covariance (including the spread of the mean term) should match the original covariance. A benefit of this approach is that it is easy to implement, however, it leads to unnecessarily fast growth of hypotheses to consider.

In the latter case, observations are allowed to be associated with tracks on adjacent discrete states, if there is a significant probability of leakage. Associated this way, the track is split into two identical ones splitting the probability as above, and where the one is projected back at the

extension of the adjacent segment (*i.e.*, as if a transition had taken place but  $x_k$  had not yet been adjusted) and then that track is updated as if it had been on the segment. This results in fewer hypotheses, as splitting is only performed when needed, but the logic is slightly more complex to implement.

### Track representation

As should be clear from the presentation above, a specific estimated track,  $i$ , is fully described by the sequence of discrete state transitions assumed to take place,  $\delta_{1:k}^i$ , and the observations associated with the track,  $\theta_{1:l}^i$ , together denoted  $\tau_{k|l}^i = (\delta_{1:k}^i, \theta_{1:l}^i)$ ,  $k \geq l$  for the state at time  $k$  with measurements from up until time  $l$ . And several different  $\tau_{k|l}$  could describe the same underlying target, with different assumptions about the discrete state sequence and what observations originate from the target.

New tracks are started where available measurements are not associated with any other existing track. The observation is used to set the general position determined from the map and the position on the road segment such that given  $Z_k$ ,

$$x_k = h^{-1}(z_k) \quad (3.8a)$$

$$\delta_k = z_k^{\delta}(\delta') = \Pr(\delta_k = \delta'). \quad (3.8b)$$

If the state  $x_k$  cannot be uniquely determined from  $z_k$  ( $h$  is not invertible), as much as possible is derived from  $z_k$  and remaining parts of the state vector is set to default values. As an example, assume  $x_k$  state with position and velocity and measurements  $z_k$  providing (possibly via a conversion) the position; then the position part of  $x_k$  is given by  $z_k$  and the speed is set to a nominal value, *e.g.* 0. This represents the possibility that the observation corresponds to a new target that has entered or newly appeared in the network.

### 3.4.2 Track hypotheses and scores

A track comprises the different hypotheses about a target, as a result of different  $\tau_{k|l}$  assumptions as described above, and a score for each of these. The different track hypotheses assumed to represent the same underlying target can be represented with a track tree, where the tree branches are a result of discrete state transitions and how the observations are associated. The score indicates the importance of the specific hypothesis. For track hypothesis  $i$  this is described as the log-likelihood ratio (Bar-Shalom et al., 2007)

$$\ell_{k|l}^i = \log \frac{p(\text{target exist} | \tau_{k|l}^i, Z_{1:l})}{p(\text{no target} | \tau_{k|l}^i, Z_{1:l})}. \quad (3.9)$$

The score can be updated recursively, as decision points are reached in the track tree.

Similar to the track state, the track score update can be divided into two steps. When the track is propagated in time, except for compensating for the probability to survive affecting all tracks alike, only the changes of discrete states affect the score when the score is divided amongst the potential transitions. The score of track hypothesis  $i$  becomes

$$\ell_{k|k-1}^i = \ell_{k-1|k-1}^{p(i)} + \log(P_s p(\delta_k | \delta_{k-1}, \hat{x}_{k|k-1}^i)), \quad (3.10)$$

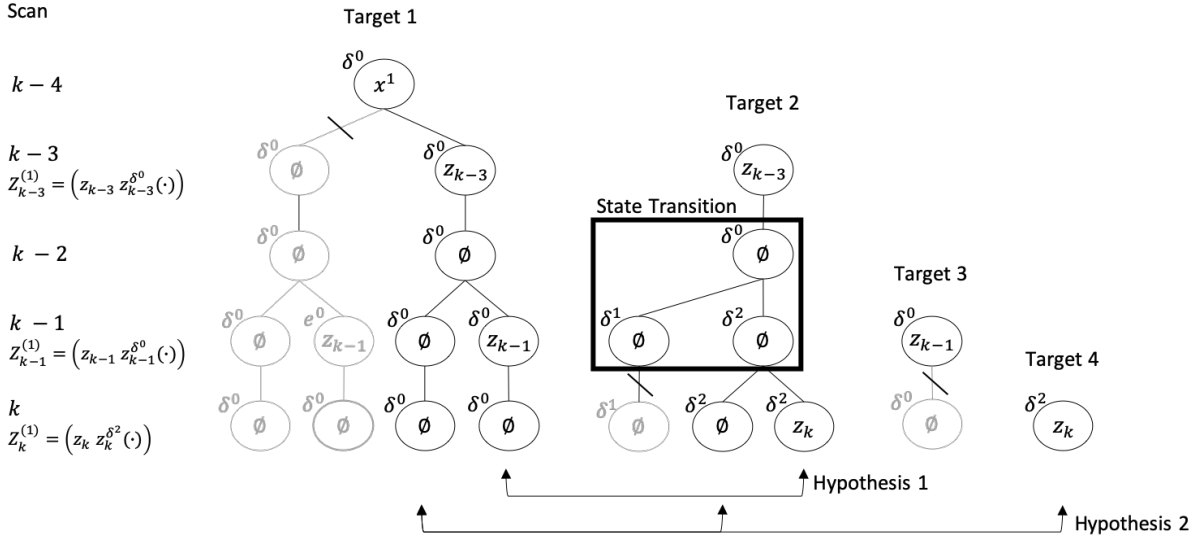


Figure 3.3: Illustration of track hypotheses and global hypotheses. A target tree can be constructed to store the hypothetical tracks of the corresponding target. A global hypothesis represents a set of all possible tracks that are compatible from different targets. The state transition occurs when  $x_{k+1}$  indicates that the target reached the end road segment, here indicated by the state  $\delta^0$ .

where  $p(i)$  is the parent hypothesis (i.e., a single prior hypotheses) of hypothesis  $i$ . It should be noted that, when no state transition is possible  $\log(p(\delta_k|\delta_{k-1}, x_{k|k-1}^i)) = \log(P_S)$ . If  $P_S = 1$  this means the score remains unchanged.

The update of the score as the result of assigning a new observation (or not) as the result of obtaining a scan is a bit more involved,

$$\ell_{k|k}^i = \ell_{k|k-1}^{p(i)} + \begin{cases} \log(1 - P_D), & \text{if } \theta_k^i = \emptyset \\ \log\left(\frac{P_D p(Z_k|Z_{1:k-1}, \tau_{k|k}^i)}{\lambda_{FA}}\right), & \text{if } \theta_k^i \neq \emptyset \end{cases}, \quad (3.11)$$

where

$$\begin{aligned} p(Z_k|Z_{1:k-1}, \tau_{k|k}^i) &= p(z_k|Z_{1:k-1}, \tau_{k|k}^i) \cdot z_k^\delta(\delta_k) \\ &= \mathcal{N}(z_k|\hat{z}_k^i, S_k^i) \cdot z_k^\delta(\delta_k). \end{aligned} \quad (3.12)$$

There are two different cases: an observation is obtained or not.

If no observation is obtained,  $\theta_k^i = \emptyset$ , the only change to the score is that the probability of the target is scaled by the probability to not observe the target,  $(1 - P_D)$ . That is, the score drops if the target is expected to have been observed.

If an observation is obtained,  $\theta_k^i \neq \emptyset$ , the probability of the target is compensated by the probability to observe the target,  $P_D$ , and the probability that the target would produce the obtained observation (which can be computed from the innovation in the Kalman filter). The no target probability is scaled with the probability of obtaining the observation without a target present,  $\lambda_{FA}$ .

**Note:** The absence of observations, sometimes denoted negative information, represents a valuable source of information as it lowers the score for track hypotheses where the target was



not observed as expected. Hence, though the track state does not change, it is important to update the track scores based on this negative information. This helps to reduce the number of track hypotheses that should be considered.

New tracks are created from observations and given the initial score  $\ell_{k|k}^i = \log(\gamma_{\text{NT}})$ , where  $\gamma_{\text{NT}}$  is used to get the right initial likelihood ratio. Terminating tracks, on the other hand, are simply dropped by cause of being too unlikely.

Given the track score  $\ell_{k|l}^i$ , the probability of track  $i$  existing can be computed as

$$p(\text{target } i \text{ exists} | Z_{1:l}, \tau_{k|l}^i) = \frac{e^{\ell_{k|l}^i}}{1 + e^{\ell_{k|l}^i}}, \quad (3.13)$$

where it has been utilized that the probability of existence and non-existence must sum to 1.

### 3.4.3 Global hypotheses

A global tracking hypothesis,  $T_{1:k}$ , is a description of a complete solution to the multi-target tracking problem; *i.e.*, the number of targets and their respective states. A global hypothesis assigns each observation to a track, or assumes it is a false observation, and determines the sequence of discrete states of all tracks. This can also be described as picking out which track hypotheses to use. The number of track hypotheses grows rapidly over time, as ideally all possible discrete state transitions and measurement associations should be explored. However, only a few of these track hypotheses are consistent with each other, fulfilling Assumption A3 which states that a measurement can only originate from one target, and can form global hypotheses. A global hypothesis, therefore, contains at most one hypothetical track from each track tree. Fig. 3.3 illustrates track trees, track hypotheses and global hypotheses, and how they are updated as new scans of observations arrive.

The probability of a global hypothesis is obtained as the product of the probability of each included track hypothesis and the probability of all unassociated measurements being false alarms. The log-probability of the global hypothesis  $l$  can be computed as the sum of the score of the contained track hypothesis compensated with a term taking false observations into consideration

$$\log(p(T_{1:k}^l | Z_{1:k})) = \sum_{i \in \text{tracks}(l)} \ell_{k|k}^i + |Z_{1:k}| \log(\lambda_{\text{FA}}), \quad (3.14)$$

where  $\text{tracks}(l)$  is the set of track hypotheses in global hypothesis  $l$ , and  $|Z_{1:k}|$  the total number of observations.

The global track score can be updated recursively, using the recursive property of the track hypotheses scores

$$\log(p(T_{1:k}^l | Z_{1:k})) = \log(p(T_{1:k-1}^l | Z_{1:k-1})) + \sum_{i \in \text{tracks}(l)} \ell_k^i + |Z_k| \log(\lambda_{\text{FA}}), \quad (3.15)$$

where  $\ell_k^i = (\ell_{k|k}^i - \ell_{k-1|k-1}^{p(i)})$  is the score increment in the current step for track hypothesis  $i$ .

In this formulation, the network constraints are naturally handled, as they are contained in the track scores, and do not require any special handling.

### 3.4.4 Complexity saving methods

To achieve a computationally tractable solution, it is necessary to lower the number of considered global hypotheses. This is done by removing too improbable hypotheses, and if possible only generate relevant hypotheses. The methods used for this are discussed in this section.

#### Generating the $k$ -best hypotheses

The number of global hypotheses grows exponentially with time, and needs to be handled to achieve computational tractability. One way to limit the complexity is to only generate a subset of the possible hypotheses. That is, to construct a subset of data associations  $\tilde{\Theta}_k \in \Theta_k$  possible in each step, such that  $|\tilde{\Theta}_k| \ll |\Theta_k|$ , which will generate the best scores. This can be posed as an assignment problem and solved with standard combinatorial optimization algorithms. Many algorithms exist to efficiently solve the optimal assignment problem based on the additive increments in the track scores, *e.g.*, the Hungarian algorithm (Kuhn, 1955), the Auction algorithm (Bertsekas, 1988), and the Jonker-Volgenant-Castanon (Jonker & Volgenant, 1987) algorithm. If combined with Murty's algorithm (Murty, 1968), the  $k$ -best assignments can be found in polynomial time. The remaining hypotheses can then simply be left unconsidered.

#### Gating

Gating is a technique to simplify the association problem, making it less computationally demanding by removing potential associations that are too unlikely. This is achieved by ignoring track to observation associations, if the observation lies outside the gate of the track.

The gate is set such that the probability of rejecting a true association is low. Using ellipsoidal gating (Collins & Uhlmann, 1992) the gate is defined by

$$D_G = (z_k - \hat{z}_{k|k-1})^T S_k^{-1} (z_k - \hat{z}_{k|k-1}) \leq \gamma_G, \quad (3.16)$$

which defines an ellipsoid around the predicted observation in which the observation must be considered. Assuming Gaussian distributions,  $D_G \sim \chi^2(n_z)$ , it is possible to design  $\gamma_G$  for a given probability to reject a correct association.

With the hybrid state space, as considered here, gating is performed in two steps, first observations where the discrete state does not match are dropped, and then normal gating is performed. Fig. 3.4 exemplifies the gating procedure, which is simplified by the fact that the continuous part of the observation is scalar such that

$$|z_k - \hat{z}_{k|k-1}| \leq \kappa \sigma_{k|k-1}, \quad (3.17)$$

where  $\kappa$  is a factor times standard deviations,  $\sigma_{k|k-1}$  in the considered dimension.

#### Clustering

Clustering is a method in which tracks that share observations are put in clusters, and then each cluster is treated independently. This is beneficial as association and track hypotheses




Figure 3.4: Road-constrained gating.

generation scales poorly with the number of involved targets and observations. When producing the final global hypotheses, hypotheses from the different clusters can be combined freely. The clustering is performed based on which tracks allow the same observation in their gate.

Clusters can be formed recursively over time. In that case, two clusters that gate in the same observation must be merged. Similarly, clusters should regularly be examined to see if they can be split, as a result of removing track hypotheses that have connected sub-clusters. For more details and implementation details the reader is referred to (Reid, 1979; Kurien, 1990; Olofsson et al., 2017). Once again, the hybrid nature of the suggested model can help simplify the handling.

### Pruning strategies: track, hypothesis and target management

Tracks are deleted based on the track score (Blackman, 2004), or  $N$ -scan sliding window (Kim et al., 2015). The first method prunes low probability hypotheses and tracks.  $N$ -scan pruning, on the other hand, traces back to the node at scan  $k - N$  and deletes the subtrees that diverge from the selected branch at that node. This represents the possibility that a target has disappeared from the scene. Furthermore, as indicated in Section 3.4.1, the network constraints cause tracks that can not be propagated to a new discrete state to be removed from the set of existing tracks. This represents the possibility that a target has exited the network. Targets are also removed accordingly, that is as soon as it has no remaining track in any hypothesis. Because the track score contain all relevant statistical information, remaining tracks lose no information during the pruning process.

## 3.5 Simulation study

In this section, the results from an empirical evaluation on three simulated scenarios are presented. The scenarios are chosen to highlight properties of the proposed NC-MHT on track-handling at intersections, data association, and measurement dependency. To evaluate the proposed approach, the NC-MHT filter was integrated into the standard (unconstrained) MHT filter from (Olofsson et al., 2017), to which it is compared<sup>2</sup> in terms of tracking performance. Note that it is difficult to provide any accurate complexity analysis, and that the proposed method has low maintenance to keep branching tracks alive, that is, carries more hypotheses that can be easily handled. The comparison is performed in terms of summarizing statistics of the trackers and

<sup>2</sup>The compared trackers are implemented in Python and C++, and the simulations were run on a laptop with a 2.8 GHz Quad-Core Intel Core i7 and 16 GB memory.

the MTT measure *generalized optimal sub-pattern assignment* (GOSPA) metric (Rahmathullah et al., 2017),

$$d_p^{(c,2)}(\mathbf{x}, \hat{\mathbf{x}}) = \left( \min_{\gamma \in \Gamma} \sum_{(i,j) \in \gamma} d(x_i, \hat{x}_j)^p + \frac{c^p}{2} (|\mathbf{x}| + |\hat{\mathbf{x}}| - 2|\gamma|) \right)^{\frac{1}{p}} \quad (3.18)$$

where  $\hat{\mathbf{x}} = \{\hat{x}_1, \dots, \hat{x}_{|\hat{\mathbf{x}}|}\}$  is a finite subset of state estimates of the ground truth set  $\mathbf{x} = \{x_1, \dots, x_{|\mathbf{x}|}\}$ , and  $\gamma$  represents the assignment sets between these two sets. Let  $d^c(x, \hat{x}) = \min(\|x - \hat{x}\|, c)$  denote a cutoff metric for any  $x, \hat{x} \in \mathbb{R}^n$ . Further, let  $\Gamma$  denote the set of all possible assignment sets. GOSPA allows a decomposition of the error into three parts: 1) localisation error, 2) missed targets, and 3) false targets. In the simulations, the GOSPA is computed based on the 2D positions of the tracks, and the parameters are set to  $c = 8$ , corresponding to a maximum error appropriate for this setting, and  $p = 2$ , corresponding to using the 2-norm which is a standard choice.

Note that GOSPA is a single unified performance metric, whereas the other quantities are presented because they represent properties that are important in MTT.

### 3.5.1 Simulation setup

The kinematic state at time  $k$  is described by the target's position and velocity. For simplicity one-dimensional motion is considered for the NC-MHT such that  $x_k = [y_k, \dot{y}_k]^T$  contains position and velocity. Nonlinear measurement models are more adequate when using real-world data, as not only position and velocity of a target is measured, but also range ( $r$ ), bearing ( $\alpha$ ), and angle rate ( $\dot{r}$ ) to a target's position can be measured (*e.g.*, Radar). The choice of an appropriate model depends on the considered application, yet this implementation uses linear-Gaussian dynamics and observations as handling the nonlinear case is a trivial extension following (3.6)–(3.7). The kinematic state motion model follows linear constant velocity (CV) given parameters


$$F = \begin{bmatrix} 1 & \Delta_T \\ 0 & 1 \end{bmatrix}, \quad Q = q^2 \cdot \begin{bmatrix} \Delta_T^3/3 & \Delta_T^2/2 \\ \Delta_T^2/2 & \Delta_T \end{bmatrix}$$

with sampling time  $\Delta_T = 1$  s, and where  $q = 0.1$  m is the process noise variance. In the simulation, targets move with initial mean speed 1.415 m/s (standard deviation is 0.215 m/s). Nonlinear effects in the target dynamics at intersections are ignored in this simulation, but can easily be considered using models describing interactions between, *e.g.*, targets at higher density junctions. For simplicity the probability of target survival is set to  $P_s = 1$ . The discrete state transition probabilities, moving targets from segment to segment, are equally distributed among the set of possible outgoing transitions from any given node.


In the simulation, sensors are mobile and independent with initial mean speed 12.3 m/s (standard deviation is 1.5 m/s), and limited field of view  $[-30, 30]$  (m) in longitudinal direction. The designed path for targets (*e.g.*, a sidewalk) is assumed all observed within a sensor's field of view. The parameters of the measurement model are

$$H = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad R = \begin{bmatrix} \sigma^2 & 0 \\ 0 & \sigma^2/4 \end{bmatrix}$$

where  $\sigma = 0.5$  m. The clutter intensity is  $\lambda_{\text{FA}}(z) = \lambda V \mathcal{U}(z)$  where  $\mathcal{U}(\cdot)$  is a uniform density over the sensor field of view,  $V = 2 \times 30$  m is the network-constrained “volume” of the field



(a) Illustration of NC-MHT and the standard two-dimensional MHT, where ground truth is represented as solid green line. The inclusion of network constraints allows for more efficient target prediction over longer time periods, and simplifies the measurement association.



(b) Tracking results at different time steps, visualizing the position and uncertainty of estimated tracks in two dimensions. Note that 2-D allows to visualize the evolution of the [NC-MHT] uncertainty, while it kept only 1-D (i.e., along the road) in the implementation. The blue crosses are sensor measurements and are at the centre of the sensor's field of view (blue circle).

Figure 3.5: Illustration and simulation results of main differences in tracking behaviour between NC-MHT (black crosses) and the standard MHT (red diamonds).

of view, and  $\lambda = 0.01 \times \text{m}^{-1}$  is the average number of clutter observations per unit volume. Further, the discrete state information included in each provided observation is accurate. In situations where several sensor platforms operate simultaneously, a sequential measurement update strategy is employed. That is, a separate measurement update step is applied for each sensor. This approach scales well with the number of sensors, and fits naturally in a distributed setting as measurements can be dealt with as they arrive.

The NC-MHT performance is compared to a standard MHT that describes the kinematic state by the target's position and velocity in two dimensions, and where the motion follows a constant velocity model. Similarly, sensor observations are given in Cartesian  $x$ - and  $y$ -coordinates.

**Note:** For the simulation, target and clutter observations were generated in the network-constrained space, yet used for both NC-MHT and standard MHT when comparing both performances. This ensures using the same measurements when comparing both filters, however, it considerably benefits the two-dimensional standard MHT. Also, the choice of linear models benefits the performance of the standard MHT, as compared to the NC-MHT.

### 3.5.2 Variable negative information

To save computations this implementation exploits the vital knowledge provided by so-called negative information, as described in Section 3.4.2. Regularly providing the filter with measurements of targets that should have been seen by the sensor, *i.e.*, non-presence, has significant impact on track likelihood scores and can help abandon unlikely tracks. Hence, the MTT should be updated with all scans, empty or not to maintain proper track scores. It is tempting to drop empty scan, as they do not affect the track estimates just the scores, in order to save time. To evaluate the importance of negative information and see its effect in the presented setting, a scheme with a varying proportion of empty scans reported to the MTT is proposed in the third scenario.

### 3.5.3 Results

Three simulated scenarios were used in this study. The first two are designed to highlight the benefits of the proposed NC-MHT as compared to the standard MHT when it comes to target prediction and measurement association. The third scenario evaluates the measurement dependency of the filter given a different number of observations. For the first two scenarios 500 Monte Carlo runs are performed, whereas the third, including two variations, is run 100 times each. The presented results are averaged over the Monte Carlo runs. For all three simulations, the maximum number of hypotheses per cluster and the minimum normalized hypothesis score are set to 50 and 4 respectively. Also note that only tracks that have been assigned at least twice with a measurement have been kept for presentation in the global hypothesis.

**Scenario 1** In the first scenario, one target initiated ahead of the fork junction was simulated for 100 time steps. True trajectories are shown in Fig. 3.6. For this scenario, 10 sensor platforms were initiated randomly distributed in the network. The scenario parameters were set to  $P_D = 0.95$ ,  $\lambda_{NT} = 0$ ,  $\lambda_{FA} = 0.6$ .




Figure 3.6: True target tracks (in different colours) for the three simulated scenarios, where black lines are road segments. In scenario 1 (left), one target is initiated on the horizontal segment before the fork junction, which eventually leads to a track split. In scenario 2 (centre), two targets are initiated well separated, before entering the two horizontal segments, potentially passing close to each other. In scenario 3 (right), three targets are initiated on random segments and more targets are born from different locations during the simulation.

This scenario particularly highlights beneficial network constraint effects for more efficient target prediction over an extended period of time. This scenario is challenging because when the target leaves the fork junction, and no measurement update is performed, correctly predicting the target state is complicated. The network constraints allow to keep track of the two different turns the target may perform as it leaves the fork intersection, whereas the standard filter predicts the target to continue straight and going off-road eventually losing track, unless an observation is obtained right after the turn. This behaviour is illustrated in Fig. 3.5a, where the standard MHT misses the turn in absence of new updates and continues its prediction in the wrong direction ( $\tau_1^1 - \tau_5^1$ ). With the arrival of new measurements, the MHT reacts by initiating a new track with along the new road segment ( $\tau_6^3 - \tau_7^3$ ).

The GOSPA performance is shown in Fig. 3.9, and summarizing statistics are shown in Fig. 3.7. A closer look at the latter unveils a higher number of tracks produced by the NC-MHT, here resulting in a higher number of clusters and global hypotheses. This is a likely bi-product of the clutter model benefiting the standard MHT and can be improved by adjusting tuning parameters. Overall, for this scenario the GOSPA results show that the NC-MHT gives smaller errors than the standard MHT filter. The GOSPA location error is larger for the NC-MCT towards the end, this is a result of the NC-MHT being better at maintaining tracks, whereas the MHT lose them effectively moving part of the location error to a missed track error. This is further substantiated by Table 3.2. Hence, in cases where maintaining track continuity is of importance the proposed NC-MHT has benefits.

**Scenario 2** In the second scenario, two targets are initiated well separated before moving onto a horizontal segment in opposite direction were simulated for 100 time steps. True trajectories are shown in Fig. 3.6. For this scenario, 20 sensor platforms were initiated randomly distributed in the network. The scenario parameters were set to  $P_D = 0.95$ ,  $\lambda_{NT} = 0$ ,  $\lambda_{FA} = 0.6$ .

This setup highlights how adding network constraints simplifies the measurement association process. This scenario is challenging for traditional approaches because when the targets are close their measurements form a single cluster, making data association difficult. Given




Figure 3.7: Statistics for simulation scenario 1

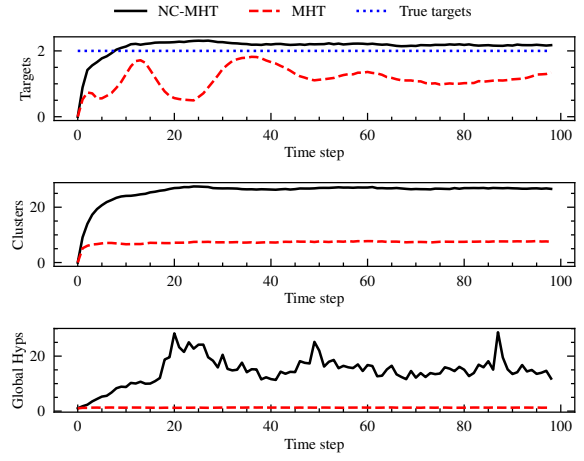


Figure 3.8: Statistics for simulation scenario 2.

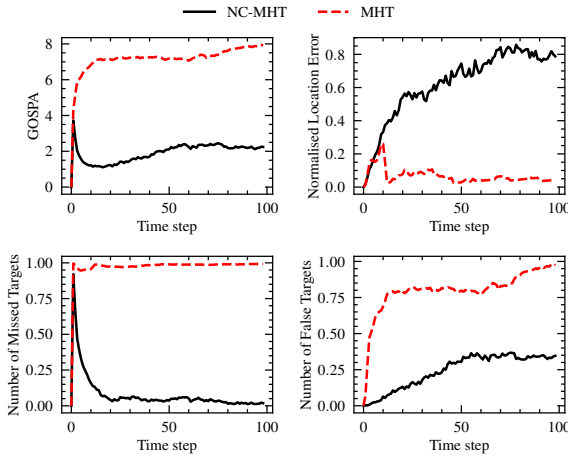


Figure 3.9: Tracking results for simulation scenario 1.




Figure 3.10: Tracking results for simulation scenario 2.

the type of targets, their speeds are assumed low and thus do not allow to distinguish their movements in two different directions that way. Fig. 3.5 illustrates how traditional approaches without network restriction need to consider all incoming measurements (if in the gate) for the measurement to track association, and possibly fail. The network constraints allow to bypass growing free-space complexity of traditional approaches and lead to more efficient gating and data association as measurements as motion is bound to a specific segment.

The GOSPA performance is shown in Fig. 3.10, and summarizing statistics are shown in Fig. 3.8. Overall, for this scenario the GOSPA results show that the NC-MHT performs better than the standard filter. The NC-MHT can keep tracks longer in time, as shown in Table 3.2. Note that, as compared to the first scenario, Scenario 2 has twice the number of sensors, hence resulting in more observations. This could explain the lower location error. As expected, results suggest that NC-MHT number of misses significantly less targets.

**Scenario 3** In the third scenario, three targets initiated at random locations in the network are simulated for 100 time steps. New targets enter in the surveillance area at different time steps



and are born at random locations in the network. Targets leaving the surveillance area is not considered here, however, can be easily handled by adjusting the probability of survival  $P_S$  for future implementations. The true trajectories are shown in Fig. 3.6. For this scenario, the performance of the filters were compared based on two realisations with different numbers of: a) number of available sensors, and b) proportion of empty scans provided to the tracker. The scenario parameters are shown in Table 3.1.

This scenario illustrates a slightly more realistic urban traffic setting and is challenging because of the many intersections, targets, and observations. This simulated scenario particularly highlights the potential of the approach for real-world applications. More precisely, Scenario 3a) highlights the (low) density of observations needed as a result of efficient predictions along the road segments in the proposed method. Scenario 3b), on the other hand, explores the importance of the negative information provided to the tracker.

Fig. 3.11 shows the GOSPA performance of the NC-MHT filter given 5, 10, 20 and 40 mobile spatially distributed sensor platforms respectively, and the performance of the MHT with 40 initiated sensors. Results show that the NC-MHT achieves a similar performance to the standard filter yet with significantly less sensors. In this simulation, the lower the number of sensors, the fewer measurements are provided to the tracker, and the less updates can be performed by the filter. This confirms results from scenario 1 suggesting that the included network structure knowledge of the NC-MHT allow for more efficient target predictions over extend period of time. The difference in track length shown in Table 3.2 supports this outcome. More sensors means more clutter measurements, but also more observations of true targets. Results also show that with more sensors, more unlikely tracks can be killed, *e.g.*, existing tracks get

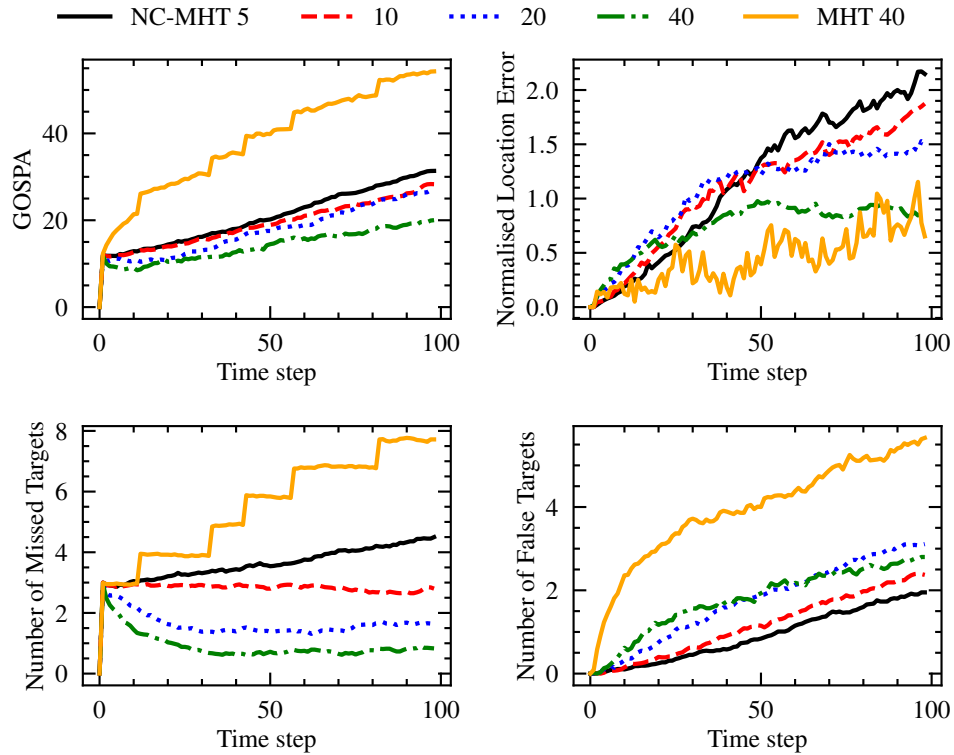


Figure 3.11: Tracking results for simulation Scenario 3a). The legend refers the number of sensors initiated.

Table 3.1: Parameters for the two realisations of simulation Scenario 3.

|      | $P_D$ | $\lambda_{NT}$ | $\lambda_{FA}$ | sensors       | empty scans         |
|------|-------|----------------|----------------|---------------|---------------------|
| S3a) | 0.95  | 2.4            | 0.6            | 5, 10, 20, 40 | 1                   |
| S3b) | 0.95  | 2.4            | 0.6            | 20            | 25%, 50%, 75%, 100% |

less likely when sensors expected a target but did not see any.

Fig. 3.12 shows the GOSPA performance of the NC-MHT filter given 25, 50, 75, and 100 percent of available sensors that provide the tracker with additional negative information, respectively, and is compared to the performance of the MHT given all sensors are included in the time update step. The results indicates that adding only a few instances of negative information can benefit the tracker significantly.

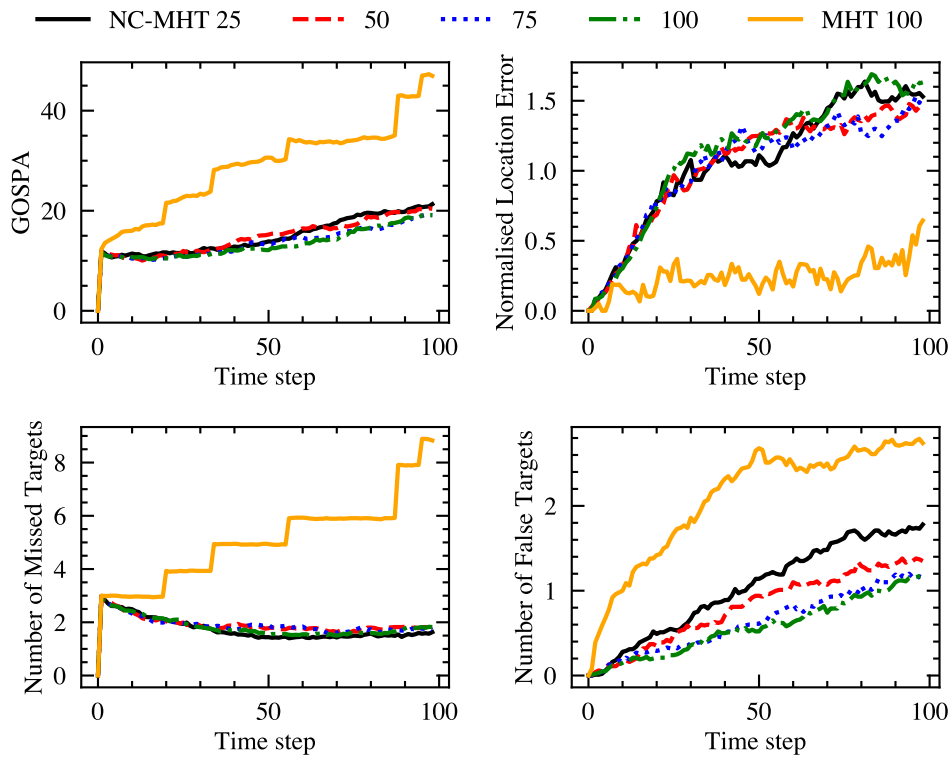


Figure 3.12: Tracking results for simulation scenario 3b). The legend refers to available proportion of available sensors (in percent) that provide the tracker with additional negative information, respectively.

Table 3.2: Results for all Scenarios 1–3.

|           | GOSPA  |        | NLE    |        | MT     |         | FT     |         | Track length |      |
|-----------|--------|--------|--------|--------|--------|---------|--------|---------|--------------|------|
|           | NC-MHT | MHT    | NC-MHT | MHT    | NC-MHT | MHT     | NC-MHT | MHT     | NC-MHT       | MHT  |
| S1        | 411.7  | 848.3  | 224.5  | 199.3  | 217.5  | 2838.8  | 2396.9 | 4732.0  | 52.8         | 26.9 |
| S2        | 159.2  | 860.0  | 37.5   | 1086.2 | 52.5   | 4532.6  | 708.7  | 2264.0  | 10.3         | 5.7  |
| S3.a - 5  | 1095.5 | n.a.   | 1313.1 | n.a.   | 9192.3 | n.a.    | 2477.4 | n.a.    | 27.1         | n.a. |
| - 10      | 1004.5 | n.a.   | 1611.6 | n.a.   | 6389.7 | n.a.    | 3072.9 | n.a.    | 33.0         | n.a. |
| - 20      | 857.3  | n.a.   | 1404.3 | n.a.   | 3927.0 | n.a.    | 3127.3 | n.a.    | 36.4         | n.a. |
| - 40      | 693.7  | 1482.5 | 894.6  | 198.8  | 1890.8 | 9758.7  | 3182.0 | 13027.5 | 39.2         | 37.1 |
| S3.b - 25 | 1055.1 | n.a.   | 1467.5 | n.a.   | 3787.2 | n.a.    | 7058.2 | n.a.    | 58.1         | n.a. |
| - 50      | 944.1  | n.a.   | 1545.1 | n.a.   | 3693.1 | n.a.    | 4827.8 | n.a.    | 49.0         | n.a. |
| - 75      | 897.5  | n.a.   | 1556.5 | n.a.   | 3836.8 | n.a.    | 3778.8 | n.a.    | 42.2         | n.a. |
| - 100     | 878.5  | 1485.4 | 1425.3 | 185.9  | 3931.8 | 11152.6 | 3421.7 | 11883.2 | 37.8         | 32.0 |

## 3.6 Outlook

In this section, we discuss technical and non-technical effects that may be induced by design decisions on the tracker, and the tracking system.

### 3.6.1 Motion and target dynamics

A comprehensive understanding of targets motion and their respective environment is an important factor for better trajectory reconstruction. Extending the proposed framework using more sophisticated kinematic models can easily be achieved, and could lead to better prediction performances on the short-term. As standard models are widely accepted for their simplicity, it seems important to keep their accuracy for trajectory estimation in balance with inherent computational cost. Yet, the longer the prediction horizon needs to be, the more important choice of the model becomes. Moreover, it is about how to make better predictions when little information is available. On the other hand, improving the scene understanding using social- and behaviour-aware models could significantly strengthen tracking and resulting trajectory estimations (*e.g.*, (Krishanth et al., 2017)), for instance, in shared spaces, when approaching signalized intersections, or during overtaking manoeuvres. This additional complexity however is a non-trivial extension to our proposed MHT approach as it breaks the assumption of independent targets. However with some approximations, *e.g.*, that the interactions do not interfere with anything else than the propagation of targets, these models could probably be used with good results and at a reasonable cost. With pedestrian and cyclist motion being influenced by weather, time of day, surrounding infrastructure and environment, as well as social and behavioural cues, future work would benefit from including such contextual information as this can lead to better long-term, and potentially network-wide, tracking results. Future extensions of this work could also substitute fixed parameters with learned distributions, which helps the tracker to better interpret measurements and tracks (*e.g.*, (Luber et al., 2011)). For instance, the rate at which new targets appear and where they appear are dependent on time and space and could similarly be learned.

### 3.6.2 Sensing infrastructures and measurements

The mobile sensing setting used in this work can conceptually be extended to any type of sensing realm where incoming measurements provide information about targets, and regularly report non-presence to improve tracking. Also since current developments hint to the integration of additional decentralized processing layers where computing power is brought to the edge, thereby fusion of additional information from additional static or participatory sensor infrastructure is conceivable and represents no extra adaptation to the proposed method.

Contextual information about the static and dynamic environment can be captured by different types of sensors. Using additional information could lead to significant improvements of the tracking system, and particularly in the data association logic, for instance as a way to strengthen the discrimination between different objects. Recent advances in deep-learning-based trackers, for instance, improve the discriminative power for visual tracking applications (Zhang et al., 2021a).

Furthermore, node transition probabilities can either be determined based on prior knowledge or set based on empiric studies about, *e.g.*, how many take a left or right turn at an intersection. Following traffic patterns this could be something that changes throughout the day and week and could be function of environment features (*e.g.*, weather, green canopy), real-time incidents (*e.g.*, events, construction), network knowledge (*e.g.*, road types, touristic routes) or infrastructure (*e.g.*, road quality).

### 3.6.3 Privacy considerations

Further improvements of the NC-MHT tracking performance can make it more suited for real-world applications, yet the gain in utility may come at the cost of the targets' privacy. The centralised scheme presented in this work, which pools all location observations to a central entity and hence may facilitate the surveillance of individuals or collectives, raises privacy concerns.

Individual location traces are akin fingerprints, as demonstrated by past research on the similarity of human mobility patterns (Alessandretti et al., 2020), their predictability (Song et al., 2010), and the uniqueness of individual traces (de Montjoye et al., 2013). But it is the realisation that with sufficient data it is possible to precisely locate, track, and possibly (re-)identify people that is a major source of concern. Standard MTT approaches have been frequently used in the location privacy community, *e.g.*, Gruteser and Hoh were among the first to use a MHT to link anonymized location samples to some individual users and recreate their traces (Gruteser & Hoh, 2005).

Anonymized location samples as considered in this paper, however, create a false sense of security because of the the spatial and temporal correlation between successive observations. Longer tracking duration, for instance, typically leads to an accumulation of information and can help infer sensitive information about individuals. The traffic density, on the other hand, can have both positive and negative impact on privacy, but is a factor beyond an individual's control. Using auxiliary data sources (*e.g.*, household address or location-based service databases) to the anonymous traces further increases the risk of inferring locations, identities, and potentially interests of an individual.

Overall, the accumulation of observations in a central database, anonymized or not, reduces the burden for attacks that can easily be automated and applied to large numbers of individuals, therefore augmenting the risk beyond individual tracking to a more collective threat.

More work is needed to study effects of the methodological, and higher-level system design space, both on target privacy and utility of future applications.

## 3.7 Conclusion

The advent of mobile sensor platforms is expected to drastically increase the amount of information collected about traffic participants (*e.g.*, pedestrians or cyclists) in populated environments, which can be of great value for multiple future applications, *e.g.*, digital twin, smart road infrastructure, autonomous driving, and advanced traffic surveillance. To make use of this information, this work presented a framework for network-constrained tracking of targets

with observations from spatially distributed and mobile sensor platforms. The key contribution represents the introduction of network bound targets into the multi-target tracking problem.

The proposed approach can be applied to any type of network-constraint environment, bypassing the growing free-space complexity of traditional MTT approaches, and use different types of observations. The generic nature of this work makes it interesting for any type of sensing setting where incoming observations provide information about targets, and potentially report non-presence of targets. Derived information can, for instance, be valuable input for traffic signal controller aiming at reducing idling times for cyclists at signalized intersections, for crowd monitoring systems measuring previously unobserved network links, or fusing observations from different sensors and road users to distribute traffic at the micro level. Overall, we see the potential to utilize this information in more places as it can be applied to more difficult problems than those addressed before.

Future work will explore the application space for next generation traffic surveillance and control systems, evaluating a range of simulated scenarios using a microscopic traffic simulator. The highlighted properties of the modified method as compared to a standard MHT can be used to predict the gains obtainable with other MTT filters. A more thorough evaluation, then focused on a particular scenario or application is an interesting direction for future research. Furthermore we motivate the need for experiments in real road traffic scenarios, which can help increase the persuasiveness of the proposed framework. It should be noted that given the (at this time) difficult access to data from a fleet of moving sensor platforms over longer time and space, any other dataset that provides information about targets and regularly report non-presence could be applied to our framework. This work hints to future directions for research in tracking and privacy engineering within the intelligent transportation realm.

## Chapter 4

# A novel data source for advanced traffic management and control

---

Previous chapters provided a comprehensive understanding of general requirements and system design, enabling the development of a method to effectively utilize data from multiple sensor nodes. This chapter leverages the insights and methods developed so far and explore specific use cases and evaluate their potential benefits and challenges.

More concretely, the objective of this chapter is to investigate and demonstrate how this new type of data can provide valuable insights and help optimize traffic control and management strategies to improve road safety and sustainability. The chapter showcases the potential benefits for cyclist prioritization in a traffic signal control setting with a comprehensive simulation. By providing a proof of concept of how this novel data source may be used in a real traffic use case, we hope to inspire various stakeholders to build upon this foundation for future research, and to think about how existing data sources could be used in similar fashion.

This chapter is published as a journal article: Vial A., Salomons A.M., Daamen W., van Arem B., Hoogendoorn-Lanser S., and Hoogendoorn S.P. (2023). Prioritizing cyclists at signalized intersections using observations from moving sensors. *Transportation Research Record: Journal of the Transportation Research Board*.

---

## 4.1 Introduction

In order to make mobility systems greener, safer, fairer, more efficient and in the end cheaper, traffic and transportation systems must continue to be geared towards multi-modal strategies. When making trips in urban environments, motorized vehicles, cyclists and pedestrians lose time when they stop and idle at signalized intersections. This increases the level of discomfort for cyclists, especially the full stops where at least one foot has to touch-down (Fajans & Curry, 2001). This is relevant as we know that the number of cyclists in cities increases with lower discomfort along the bike network (Pashkevich, 2013). The urge to reduce delays for cyclists at signalized intersections has recently drawn increasing attention from the traffic community (Wierbos et al., 2021; Claus, 2022). The motivation of this study is to investigate the potential of early cyclist detection for responsive signal control that prioritizes cyclists, so this mode of transport will suffer less from stops and long delays.

In the Netherlands, signalized intersections are controlled by *vehicle-actuated control* (VA). VA control uses information about the current traffic state to control the use of the conflict areas, *i.e.*, areas where two conflicting movements meet. The order in which directions will receive green is based on a fixed control structure, with stages of non-conflicting movements constructed with the aid of pairwise conflicts (Furth et al., 2012). The order of the phases is red-green-yellow, and the yellow time is based on the speed of the traffic and the distance where traffic can choose whether to proceed or to stop, the so-called dilemma zone. The actuation is mostly done by means of loop detector information (Reggiani et al., 2022). VA is phase control, the green duration of an individual movement is determined by the occupation of its detectors. Although multiple detectors are recommended for all traffic modes (van Dijk et al., 2018), two to three inductive loops are typically used upstream of the intersections for cars, while for cyclists, only one inductive loop is placed near the stop line.

More recently, VA control has been extended to enable certain movements to be prioritized, *i.e.*, *VA control with priority* (VAP). These extended, so-called ‘intelligent’ traffic light controllers (iTLC), are enhanced by communication with the traffic (van der Vliet et al., 2016). The aim of iTLC is to inform the approaching traffic, to optimize the control and to prioritize target groups. For example, a public transport vehicle, a truck, or an emergency vehicle can initiate the end of green times of conflicting movements to pass as soon as possible. Similarly, VAP can be used to prioritize cyclists based on conditional priority, *e.g.*, with rain sensors that are used to prioritize cyclist during rainfall only (Fietsberaad, 2012; Lusk, 2016). Eventually the iTLC will improve the estimation of traffic states (Poelman et al., 2020). The communication with cars is done via on-board computers or mobile apps, while cyclists rely on communication enabled by Radar on a fixed location (Lai, 2021), or from smartphone apps (Verbeeke, 2020). The number of cyclists using mobile apps, however, is still low as not all cyclists are willing to use an app or remember to start an app (Verbeeke, 2020). This mismatch results in a higher share of cars communicating with the iTLC than communicating cyclists.

Because cycling is encouraged in the Netherlands (Reggiani et al., 2022; Arellana et al., 2020), more solutions are needed to balance out the lack of information about cyclists approaching the intersection. Eventually, this will make it possible to prioritize this environmental friendly traffic mode more effectively than is currently the case. Other solutions to increase traffic controllers’ situational awareness of cyclists is to install additional sensing infrastructure,



*e.g.* extra detector loops, fixed cameras or radars (Claus, 2022). This approach, however, comes at a considerable cost of deployment and maintenance (Klein et al., 2006).

Recent studies utilize connected automated vehicles (CAVs) as a mean for learned eco-driving control policies to reduce idling time at signalized intersections (Jayawardana & Wu, 2022). The broad sensing, processing and communication capabilities of future CAVs make them well-suited for advanced traffic monitoring and control applications. As they must continuously sense the environment in which they operate, CAVs may also provide relevant traffic information about cyclists, and can thus be understood as mobile sensing platforms (Vial et al., 2020). Observations generated from their field of view can be shared using wireless communication with nearby vehicles (*e.g.*, for cooperative vision) or infrastructure such as traffic signal controllers (*e.g.*, with iTLC). The benefits of using observations from CAVs, is that the position and speed of cyclists may be observed by multiple CAVs at different locations (*e.g.*, to confirm the presence at a location), while cyclists do not need to start an app. Since prioritization of cyclists may increase the delay of other traffic modes, robust observations are necessary to avoid giving priority too early or for too long.

The main ambition of this paper is to show how augmenting the situational awareness of traffic signal controllers using observations from CAVs, can enable prioritizing cyclists and reduce lost times in the control cycle in an effective way. In this paper, the following contributions are made:

1. propose a control method that uses available cyclist observations from moving sensor platforms (here exemplified with CAVs) to prioritize and optimize traffic conditions for cyclists,
2. quantify the effect on cyclists and cars delays, on the number of stops and on effective green use for cyclists, for different levels of car and bicycle demand,
3. highlight to which extent different traffic and environmental conditions may influence these results (*i.e.*, settings of the CAV, penetration rate, occlusion),
4. show how observing the absence of cyclists appears to be valuable controller information (*e.g.*, reducing lost time by phase truncation)

We showcase this with a vehicle-actuated control application that prioritizes cyclists, and use a microscopic traffic simulator to simulate cyclist and car traffic approaching a signalized intersection. Note that in this work we consider cyclists to be indistinguishable from each other, that is, no additional information from sensors (*e.g.*, clothing colour) that could help strengthen the discrimination between different cyclists is assumed to be given. Available observations from CAVs are used to estimate the dynamic states (*i.e.*, positions and speeds) of an unknown and varying number of cyclists using a multiple-target tracking (MTT) approach. The detector information is then augmented with the estimated cyclist states as input for the cyclist priority request logic of the proposed *cyclist prioritising vehicle-actuated controller with tracking* (VAT). The cyclist and car delays from the implemented VAT are empirically evaluated against a state-of-the-art VA and VAP controller. This work motivates for more optimized designs that better balance delays in more involved traffic situations.

The remainder of this paper is organized as follows. The next section gives an overview of the proposed VAT approach. Then, we recall the MTT approach used to track cyclists from

individual observations and present related models used for this study. Models used for priority requests and lost time reduction in the control cycle are introduced in the fourth and fifth section. Experimental design and simulation setup are presented in the sixth section, followed by the results from an empirical evaluation. Conclusions and future work are presented at the end of the paper.

## 4.2 Approach

The objective of the proposed method is to use additional information from moving sensor platforms<sup>1</sup> to inform, prioritize and optimize intelligent traffic signal controllers. The general idea and main elements of our approach, are outlined in this section and depicted in Figure 4.1.

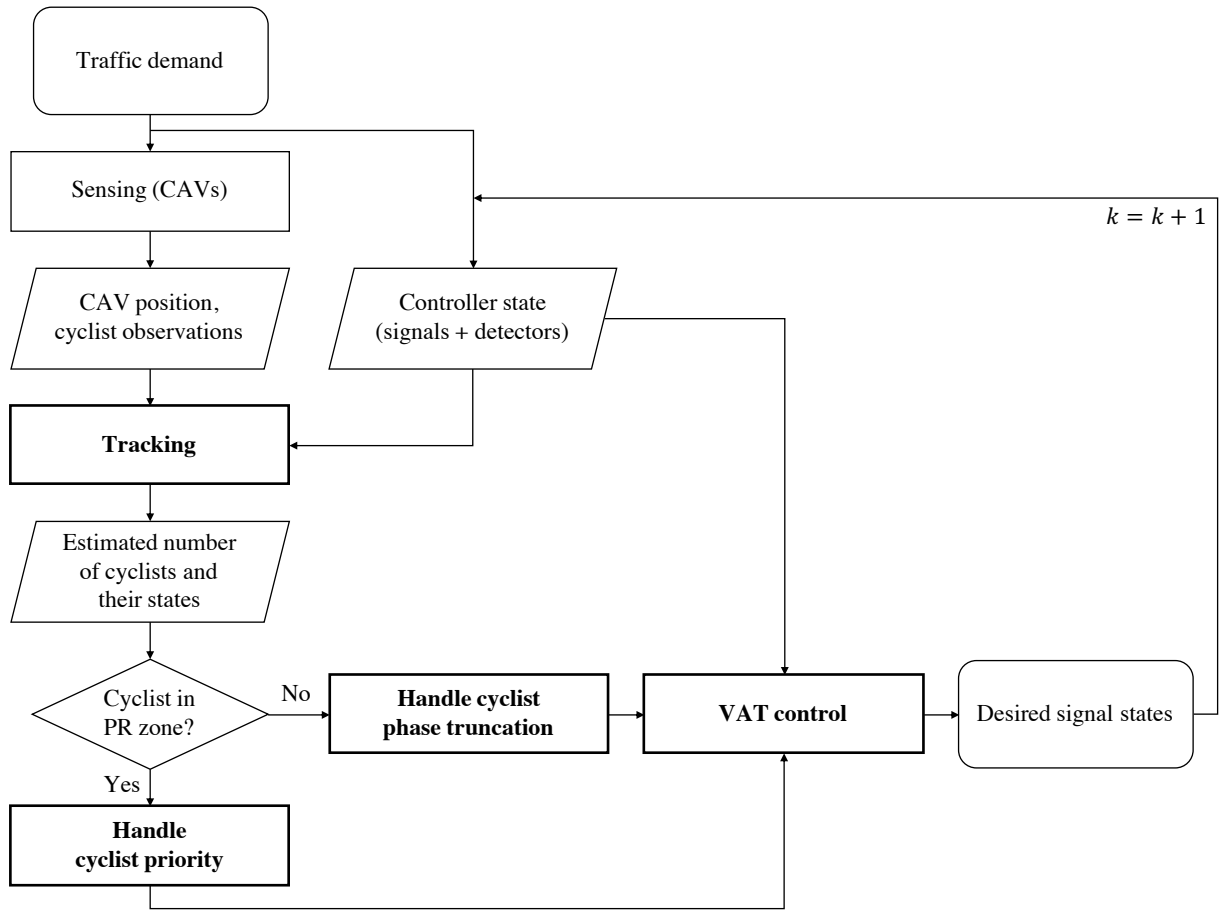


Figure 4.1: Bicycle priority Vehicle-Actuated Control with Tracking (VAT) using observations from moving sensor platforms (e.g., CAVs). PR zone: zone where priority can be requested.

When a cyclist enters the field of view (FoV) of a sensor platform, an observation may be generated. Based on the architecture of the iTLC, where road users share their position, speed and intended route by wireless communication with the traffic controller (Fløan et al., 2016),

<sup>1</sup>This article uses the term *moving sensor platform* as a more general description of technological sub-systems that generate data of their surroundings for their safe functioning (e.g., CAVs, drones, or other types of robots). Later in the paper, we will showcase the proposed approach with CAVs in simulation.

this work assumes sensor platforms to similarly share available observations of cyclists with the traffic controller. In order for the traffic controller to make use of available observations over time, the dynamic state of observed cyclists (*e.g.*, position and speed) must be estimated along the cycle path so that, at a suitable moment, the priority request of the signal controller can be activated. The way the iTLC handles the input from sensor platforms to grant priority to individual cyclists, or a platoon of cyclists, is comparable to the way data from static camera or Radar is handled <sup>2</sup>.

The unknown and varying number of cyclist, however, complicates the task of associating available observations with cyclists present on the cycle path. In addition, sensor platform observations are not perfect. They may be noisy, include clutter or miss a cyclist because of occlusion. And with no *a priori* information about which observation originates from which existing or newly detected cyclist, the many possibilities of assigning a measurement to an individual complicate the task. We use a multiple-target tracking approach to infer the number of cyclist present on the cycle path and estimate their dynamic state at each time instant.

For the case cyclists are estimated to be located in the priority request zone<sup>3</sup>, a priority request logic is initiated. The VAT controller then ends green time for conflicting traffic at the next possible moment and the desired signal state (*i.e.*, green) is set for cyclists. In this way cyclists can be prioritized earlier, unlike VAP, where the cyclists have to continue up to the stop line and have to stop before the priority will be granted. For the case cyclists are estimated not to be located in the defined priority request zone, a phase truncation logic is initiated that will reduce the lost time in the control cycle.

## 4.3 Inferring the number and states of indistinguishable cyclists

In order for the traffic signal controller to act according to information provided from sensor platforms' field of view, we use a MTT approach to infer the number of cyclists and estimate their dynamic states (*e.g.*, position, speed) from a sequence of available observations. MTT is an extension to the state estimation problem, including a data association component as it is not known which observation originates from which cyclist, not to mention which observations are clutter or missed detections. In this section, we introduce the used MTT approach, relevant state space models, and discuss the applicability and assumptions made using this approach in a traffic signal control setting.

### 4.3.1 Road network-constrained MTT

We use the network-constrained multiple hypothesis tracking (NC-MHT) approach proposed in (Vial et al., 2023a) to update the belief of a target's position along a defined road network from a temporal sequence of observations.

---

<sup>2</sup>An example where bicycle priority is applied using smart cameras detecting bicycle platoons of three or more cyclists can be found in the municipality of Hengelo (NL) (Crow, 2022), while bicycle priority using cyclists observations from static Radar was applied in the municipality of Delft (NL) (Lai, 2021).

<sup>3</sup>Area upstream of the traffic signal where the decision about cyclist priority is made.

At each time step, an association event represents a set of valid assignments of observations to existing tracks (of already observed targets), new targets or false observations, where a track corresponds to a data association sequence  $\tau_{1:k}^i$  for target  $i$  up to time  $k$ . Each track is modeled as single target tracking filters (*i.e.*, Bayesian filtering) following prediction steps or updates in case it is associated to an observation. The decision whether or not to update an existing track with an observation is made using *gating*, a technique used to simplify the association problem by removing potential associations that are too unlikely. If the observation lies outside the gating area predicted for a track (based on the state estimate, the next observation of the track is predicted), the track to observation association is ignored.

A complete solution to the MTT problem is given by a global tracking hypothesis, which represents the inferred number of targets and their respective states along the road network. We refer the interested reader to the original work (Vial et al., 2023a) for more details on, *e.g.*, data association logic, track likelihood computations or handling of the road constraints, as we focus on the target state description and estimation models for the rest of this section.

### 4.3.2 State space models

Let us assume that cyclists evolve independently of each other and the state of each cyclist at time  $k$  is represented by a vector

$$X_k = \begin{pmatrix} x_k \\ \delta_k \end{pmatrix},$$

where, the actual position of the cyclist is a combination of the general position determined from the map  $\delta_k$  (*i.e.*, cycle lane), and  $x_k$  which specifies the position, velocity and acceleration on the cycle lane, such that  $x_k = [r_k, \dot{r}_k, \ddot{r}_k]^T$ .

#### Cyclist motion model

A nonlinear dynamic model describes the motion of each cyclist such that

$$x_{k+1} = f(x_k, u_k) + v_k, \quad v_k \sim \mathcal{N}(0, Q_k), \quad (4.1)$$

with a control input  $u_k$  (*e.g.*, to specify the traffic signal state at each time step), and  $v_k$  is process noise assumed to be white and Gaussian with covariance matrix  $\text{cov } v_t = Q_k$ .

The kinematic state motion model of cyclists approaching the intersection follows linear constant velocity if not in so-called human-vision distance to the traffic light.

In human-vision distance to the traffic light, on the other hand, we assume the signal status to have an effect on the cycling behaviour. We use a simple model that describes cycling movements in response to the colour of the traffic light as proposed in (Dabiri et al., 2022). Thereby, the acceleration/deceleration of a cyclist is represented as follows,

$$\bar{r}_k = \begin{cases} -\frac{\dot{r}_k}{C_S(k)\Delta t}, & 0 < l^S - r_k < D_h, \\ & \text{if signal is red,} \\ 0, & 0 < l^S - r_k < D_h, \dot{r} > \dot{r}^d, \\ & \text{if signal is green,} \\ \dot{r}' \left( 1 - \left( \frac{\dot{r}_k}{\dot{r}^d} \right)^2 \right), & 0 < l^S - r_k < D_h, \dot{r} < \dot{r}^d, \\ & \text{if signal is green,} \end{cases} \quad (4.2)$$

with discretisation time  $\Delta t$ , comfortable speed  $\dot{r}^d$ , and comfortable acceleration of the cyclist is denoted  $\dot{r}'$  at time step  $k$  and where  $C_S(k)$  denotes the number of required time steps in order to fully stop before the intersection,

$$C_S(k) = \max \left( 1, \left\lfloor \frac{2d_j(k)}{\dot{r}_k \Delta t} \right\rfloor \right), \quad (4.3)$$

where  $d_j(k)$  denotes the distance between the cyclist and the traffic light at time step  $k$ . This model is applied according to the following cases, where  $l^S$  represents the position of the stop line and  $D_h$  represents the human-vision distance from the traffic light: In the first case, while the cyclist is in human-vision distance to the traffic light and the signal is red, the cyclist decelerates until an eventual full stop ahead of the traffic light. In a second case where the light gets green while in the human-vision distance, no acceleration is applied to the cyclist, unless the current speed is lower than the defined comfortable speed.

As for the case the cyclist has passed the intersection, the kinematic state motion model follows constant velocity if cycling at comfortable speed, otherwise the acceleration is represented according to the last case of Equation (4.2).

### Observation model

Let us assume moving sensor platforms to have a limited FoV, a known position along the road network and that observations may be generated at any time. Each observation originates from a single cyclist, or is a false detection. When a cyclist is visible to a sensor platform at time  $k$ , and thus in its FoV, the cyclist is detected with probability  $P_D$ , where the available observation provides a measurement about the cyclist's general position from the road map (*e.g.*, the probability of being on a specific cycle lane) and specifies the position and speed along that road segment (*e.g.*, where on the cycle lane). A single sensor platform can observe multiple cyclists at the same time and thus provide multiple observations at one time step, yet each target produces at most one observation. Given multiple sensor platforms operating, a single cyclist may also be observed by several sensor platforms simultaneously.

Note, that the absence of observations represents a valuable source of information as it lowers the score for track hypotheses where the target was not observed as expected. This helps to reduce the number of track hypotheses that should be considered.

### 4.3.3 State prediction and update

As mentioned earlier, tracks are estimated using Bayesian filtering following prediction steps or updates in case it is associated to an observation. We use an extended Kalman filter (EKF) (Jazwinski, 1970) in case of nonlinear motion, and the standard Kalman filter (KF) (Kalman, 1960) for the linear case. When an observation is obtained, the continuous part of the cyclist state can be updated, provided the discrete state is possible in the observation. If no observation is available, the update step is simply skipped, and if several are available, the step is repeated. We refer the interested reader to (Vial et al., 2023a) for more details on network-constrained effects and assumptions in the context of the NC-MHT framework.

### 4.3.4 Applicability and assumptions in a traffic signal control setting

The flexibility of the used NC-MHT framework, particularly the formulation of the network-constrained target state, makes it generally suitable to associate available measurements with cyclists present on the cycle path. As compared to more standard MTT approaches in free space, the inclusion of network constraints allows for more efficient target predictions over extended periods of time and simplifies the measurement association process. As this work considers moving sensor platforms as new source of cyclist information, it is reasonable to assume the knowledge about the network structure and the sensor platform position always to be known.

Cyclists overtaking other cyclists or approaching a signalized intersection can influence one another. Their motion can, for instance, be affected by the behaviour of others, the surrounding environment or their destination choice. More advanced motion models, including interactions of targets have the potential to be utilized to improve tracking, improving predictions in particular over extended periods of times without observations. However, interactions between targets strictly speaking violate a key assumption of the used framework, and more generally of all classic tracking frameworks, as they treat targets independently to save complexity. Extending the method to take these interactions into consideration is non-trivial and more importantly has a very negative impact on computational complexity. To prove the concept of our application, we therefore consider cyclists to move independently as bypassing the violation of this basic assumption and using more involved behavioural models is out of scope for this study.

At the same time, cyclists respond to the colour of the traffic light, *e.g.*, by adapting their speed. In contrast to the interactions between cyclist, the traffic signal colour is a fixed input at every time step and can thus considered in this study as additional control input for the cyclist motion in the used approach.

The used NC-MHT framework models transitions in a road network as a discrete Markov Chain, where a target transitions from one road segment (*e.g.*, cycle lane) to another once it has reached the end of that segment (*i.e.*, the intersection), according to a computed transition probability. When applied to a traffic signal control setting, *e.g.*, where cyclist approach an intersection, the provided network representation allows to specify different cycle lanes. Each sensor platform observation could thus include the probability that the cyclist was observed on that cycle lane. To limit computational complexity, this paper uses a simplification where the cycle lane is assumed to be accurately provided by the sensor platform (*i.e.*, only one cycle path has non-zero probability).

Note that for some traffic signal control applications (as exemplified later in this study) the complete intersection area (*e.g.*, conflict areas) must be included in the continuous description of the cyclist motion state, so the network representation used in the NC-MHT must be such that cyclist can be tracked not only before and after an intersection, but also at the intersection. Ignoring the presence of a cyclist in such crucial zones would have considerable negative effects for overall safety and efficiency of the controller.

## 4.4 Arrival time and priority request

To give priority to cyclists at the right moment, the arrival time of the cyclist must be calculated. Giving cyclists priority will have an effect on the delay of other traffic, and to avoid causing excessive delay to conflicting vehicle movements, the priority of the bicycle movement will be conditional.

### 4.4.1 Arrival time calculation

To find the optimum moment to give a cyclists priority, the position  $r_k$  and speed  $\dot{r}_k$  are extracted from the tracks included in the best global hypothesis of the tracker. The time step between new estimated trajectories, between  $k$  and  $k + 1$ , is  $\Delta k$ . The speed of the cyclists is assumed to be constant between the two observations. After  $n$  time step  $t_s$  the position of the cyclists is determined by  $r_{k,n} = r_k + \dot{r}_k n t_s$ , where  $n \leq \Delta k / t_s$ . The time it takes a cyclist to arrive at the stop line is:

$$t_{arr} = \frac{r_{st} - r_{k,n}}{\dot{r}_k}, \quad (4.4)$$

where  $l^S - r_{k,n}$  is the distance the cyclists has to cover at time step  $k + n t_s$  to the stop line  $l^S$ .

### 4.4.2 Conditional priority request

The priority start of the bicycle movement depends on the arrival time  $t_{arr}$  of the cyclists, on the controller state, and on the priority conditions. First we discuss the priority start moment, followed by the conditions for which priority is given, or extended.

When the conditions favour priority, the start of the priority request depends on the controller state:

1. During green, yellow or clearance times of parallel movements the priority is requested at the moment the cyclist is about to brake, this is  $t_{Amin}$  seconds before arriving at the stop line. The light can turn green immediately after requesting:

$$t_{arr} \leq t_{Amin}. \quad (4.5)$$

2. During green, yellow or the clearance time of conflicting movements: also here the signal needs to turn green on the moment the cyclist starts to brake. However, if a conflicting

signal  $i$  is green, it takes time to end this signal, so the remaining time  $t_{r,i}$  before the bicycle green can start should be taken into account. The request is done if:

$$t_{arr} \leq t_{Amin} + \max_i(t_{r,i}), \quad (4.6)$$

The remaining time depends on the yellow time  $L_i$  of the conflicting movement  $i$ , the clearance time  $C_i$  between the conflicting movement and the bicycle movement, and  $G_{end,i}$ , the time that has passed since the green time of movement  $i$  ended:

$$t_{r,i} = L_i + C_i - G_{end,i}, \quad (4.7)$$

The conditions for starting the priority depend on the green duration of conflicting movements, and on the number of times priority was realised for the bicycle movement:

1. do not truncate conflicting movements too soon, at least  $p_c\%$  of the maximum green should be given to a conflicting direction,
2. limit the number of priority realisations, after the bicycle priority ends, it should not start again before the lost time has expired. This ensures that one of the conflicting movements that have a green request, will receive green.

Also priority extension should be limited, do not extend the green time for cyclists too long:

1. the green may extend at maximum  $p_b\%$  of the maximum green,
2. extend the priority only if a cyclist will arrive at the stop line within  $t_{Amax}$  seconds after the predecessor moved over the stop line.

Since VAT priority is granted when a cyclist is observed and there is a conditional extension of the green time, a single cyclist that arrives will be prioritized over a platoon of cyclists that will arrive after the  $t_{Arr}$ , or after the maximum green extension of the bicycle green has been reached. This means that the VAT controller does not optimise the green time with respect to the arrival of platoons, nor does it consider platoon dispersion.

## 4.5 Reducing lost time by phase truncation

Moving sensor platform observations can be used to prioritize approaching cyclists, yet can also be used to optimize the control by reducing unused green and lost time in the control cycle by truncation of the minimum green time, the yellow time or clearance time, here indicated as “phase truncation”. Since two separate areas need to be observed, the area from the priority request location to the stop line (PR zone), and on the intersection (clearance zone), the phase truncation is done in two steps. First the minimum green and yellow truncation is considered related to the PR zone, secondly, if the conflict zone is clear, the clearance time is ended.

If it is observed that there are no cyclists present at the PR zone, and the bicycle signal is showing green because the minimum green time has not expired yet, the minimum green time can be ended. The gap-out time of the detectors is not needed, since it is certain the queue is



cleared. If there are no cyclists in the PR zone (which is upstream of the dilemma zone), no cyclist can decide to go when yellow is shown, so the yellow time does not have any function and can be ended. This truncation should take into account the speed of the fastest cyclists  $\dot{r}_{c,fast}$  (e.g., e-bikes) and the duration of the yellow signal  $L_c$  :

$$d_{trunc} = \dot{r}_{c,fast} L_c \quad (4.8)$$

If cyclists are not present in the priority request location and in the conflict area, the clearance time can be ended, the conflict area is clear for conflicting movements. The safety risks of phase truncation is reduced by the demand that multiple sensor platforms should observe these areas, the priority request location and the conflict area, and these sensor platforms should provide redundant observations that there are no cyclists present.

## 4.6 Experimental design and simulation setup

The aim of this case study is to quantify the potential of using cyclist observations from moving sensor platforms, here exemplified in a CAV use case, with the implemented VAT approach that prioritizes cyclists by tracking, by answering the following questions:

- Q1:** How does the VAT, VA (no priority) and VAP control (priority on detection) compare in terms of car and bicycle delay and stops, for different levels of car and bicycle demands?
- Q2:** What is the effect of using VAT on the efficient use of green for cyclists?
- Q3:** How does the penetration rate and the FoV of CAVs affect the cyclist delay?
- Q4:** What is the value of observing the absence of cyclists?

This section first provides more details about the scenarios simulated in this study. We then describe the designed case study, and eventually elaborate on assumptions and design choices made for the implementation of the numerical experiments.

### 4.6.1 Simulation scenarios

To answer the questions listed above, we evaluate the benefits of the proposed VAT approach with five simulated scenarios using PTV Vissim 11 microscopic traffic simulator (PTV, 2016). The road network is synthetic, that is, not based on an existing intersection and thus the microsimulation network is not calibrated with existing data. The network is made sufficiently complex to show the proof of concept, yet simple enough to make the set-up feasible. The demand levels are based on realistic demands (Godefrooij, 2016).

- To answer **Q1**, we evaluate the benefits of the proposed VAT approach against VA and VAP implementations in terms of the effect of the bicycle demand (100-800 c/h), for various levels of car demand (50%-150% of the baseline car demand, given in Figure 4.2).

- For **Q2** we assess how VA, VAP and VAT handle green actuation and duration for the different bicycle demand.
- For **Q3**, different variables of the tracker were changed to compare different VAT scenarios. The VAT base scenario was chosen with 300 c/h, 50% CAV and 30m FoV, as well as no occlusion or phase truncation. The baseline car demand is given in Figure 4.2. The probability of detection for each CAV is set to  $P_D = 0.95$ . The penetration rate of CAVs was evaluated based on 50%, 10% and 5% CAVs. Note that, the baseline case has 50% CAVs since results from preliminary tests did not show much variation with a higher penetration rate. The field of view was set to 30m, 20m and 10m.
- Question **Q4** can be split in two situations. First, if a CAV does not observe any cyclists, this can be the consequence of occlusion, thus occlusion events are added and compared to the VAT base scenario described above. Second, if observations from multiple CAVs provide redundant information about the absence of cyclists and no occlusion takes place, we can evaluate the benefits of the proposed VAT approach against VA and VAP implementations in terms of phase truncation. In particular we are interested in minimum green, gap time, yellow time and clearance time truncation.

Given stochastic effects and priority, we used 7 different seeds to simulate, with a simulation duration of 900 s (10-15 cycles), with sampling time  $\Delta T = 1$ s for the tracker, and the control is updated every  $\Delta T = 0.1$ s yet assuming a constant speed during that second.

#### 4.6.2 Case study

We simulate cyclists and vehicles moving along the road network depicted in Figure 4.2. This intersection is close to a realistic situation and sufficiently complex in seeing the effect of prioritization. For simplicity only CAVs parallel to the cycle lane can observe cyclists. To mimic a realistic control, a three block-control was implemented, but with only one direction that is not parallel to the cycle lane. Cyclists travel from South to North on the (red) cycle path. A loop detector for cyclists is placed at the stop line, as is most commonly used in the Netherlands (Reggiani et al., 2022) (a benchmark country in terms of bicycle infrastructure). As mentioned in the introduction of this paper, the detector is used as actuator for the VA-, VAP- and VAT-approach. Vehicles present on the network are either conventional non-sensing (“car”) or equipped with sensing capabilities allowing them to generate observations about objects in their field of view (“CAV”). CAVs that have the cycle path within the defined FoV can generate cyclist observations at any time. The cyclists moving along the cycle path are either ‘common’ cyclists with an average speed of 15 km/h, or fast cyclists (*e.g.*, e-bikes) with a speed of 25 km/h. These speed differences are introduced to challenge the arrival time estimation.

#### 4.6.3 Network representation

According to the tracking framework introduced earlier, all observations received from CAVs are network-constrained, which can theoretically be represented as a Markov model in the form of a directed graph. As indicated earlier the complete intersection must be included in the continuous description of the cyclist motion state. We thus represent the cycle path entirely as a

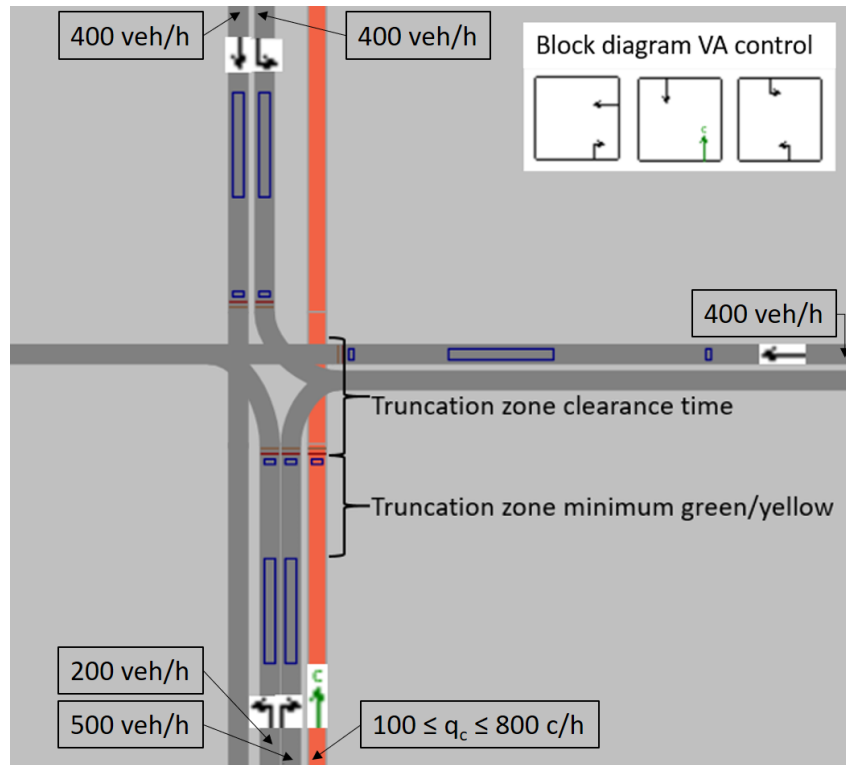



Figure 4.2: Intersection used in microsimulation, with the control block diagram. The bicycle movement is depicted by the arrow with 'c'. The demand for all movements is given in the figure.

single line segment. This design choice allows us to have cyclist observations and propagation at the complete intersection area, without issues of cyclists transitioning from one segment to the next.

#### 4.6.4 Field of view and occlusion

As indicated earlier, the size of the FoV determines if a CAV can observe the cycle path, and if so, which part. Only cyclists within the FoV can be observed by a CAV. The shape of 2D-FoV can be approximated best by an ellipse (D. Göhring & Ganjineh, 2011), but for simplicity, we approximate the FoV by a square. Figure 4.3 shows that this approximation is reasonable with respect to the actual size of the FoV. For the baseline VAT the size of the FoV is set to 30m. The size of the FoV is determined by the lateral distance to the cycle path. CAVs with a FoV of 30m can observe the cycle path being three lanes away from the cycle path, while with a FoV of 10m, only CAVs on the lane directly next to the cycle path can observe cyclists. In more practical terms, the larger the FoV, the more information may be shared with the tracker, which leads to higher certainty in cyclist estimates.

For occlusion events, we implement a simple model where only lines of sight between the CAV and the cycle path, perpendicular to the cycle path are considered. If another vehicle is partly within the perpendicular line of sight, this part of the cycle path cannot be observed by the CAV. Figure 4.3 illustrates the simplified occlusion model. Close to the stop line waiting vehicles can obstruct this view, so the occlusion is worst near the stop line. The consequence



*Figure 4.3: Illustration of the simplified occlusion model. The observed bicycle path areas are hatched, the two cyclists within the circle are not observed because of the occlusion due to the red car.*

of this event is that if the tracker gets conflicting information, that is, from CAVs that observe the cyclists without occlusion in their FoV and from CAVs that cannot observe a cyclist but should give their position with respect to the cycle path, the tracker may (depending on track likelihood computations) end the trajectory.

#### 4.6.5 Implementation details

The Vissim network is controlled by a vehicle-actuated controller (Trafcod, (Furth & Muller, 1999)) which reacts on the occupancy of detectors. The COM-interface gives information to Controller adjustment in Matlab, Matlab passes (a part of) the information through to the Tracking algorithm in Python and C++, information from the Tracking algorithm is input for the Controller adjustment. The Controller adjustment decides when a direction will start or end green or end yellow, and based on this, certain detectors will be cleared or occupied, to make sure the signal timing is as intended.

### 4.7 Results

In this section, the results from the empirical evaluation are presented, for the proposed VAT controller and two state-of-the-art controllers, VA and VAP:

- **VA: Vehicle Actuated control**, using detectors, without any priority,

- **VAP: Vehicle Actuated** control where cyclists are **Prioritised** on the stop line detector,
- **VAT: Vehicle Actuated** with cyclist priority using CAV-based **Tracking**.

The base scenario is with 50% CAV and 30 m FoV, without occlusion or phase truncation and the baseline car demand from Figure 4.2.

Figure 4.4 shows how the tracker deals with the cyclists observations, and estimates the cyclist trajectories, by comparing to the actual trajectories of the cyclist in the simulation. This figure also highlights how a track initiated by a false observation results in a unjustified priority start of green.




Figure 4.4: Trajectories of cyclists, blue line: actual trajectory (true positions), red: observed by CAV within FoV, and the position of the cyclists according to the tracker algorithm (estimated positions). At 300m the stop line is indicated with the signal phases (green, yellow, red). In grey circles false observations are indicated.

#### 4.7.1 Q1: VAT vs. baseline control scenarios

In answering **Q1**, Figure 4.5 presents an intuition of the performance of the proposed VAT approach against the baseline cases used in this study. It presents the percentage of cyclist stops, bicycle and car delay as function of the demand. As can be seen from Figure 4.5a, VAT shows a reduced number of required stops for cyclists as compared to both VAP and VA. In particular, observing cyclists ahead of the stop line in a priority scenario with tracking (VAT) considerably improves the number of stops when compared to a priority without tracking scenario (VAP), where there is a high probability that a cyclist stops on the detector since the detector is close to the stop line.

The delay as function of cyclist demand, see Figure 4.5b, is smallest for VAT with a low demand while it is highest for VA. This result is a logical consequence given VA request can

only be honoured when the block of the cyclists movement can be activated. If VAT is compared to the VAP results, the difference between the delay of cyclists is smaller than for the stops, the effect of prioritization on delay reduction is larger than the effect of being observed further ahead of the stop line.

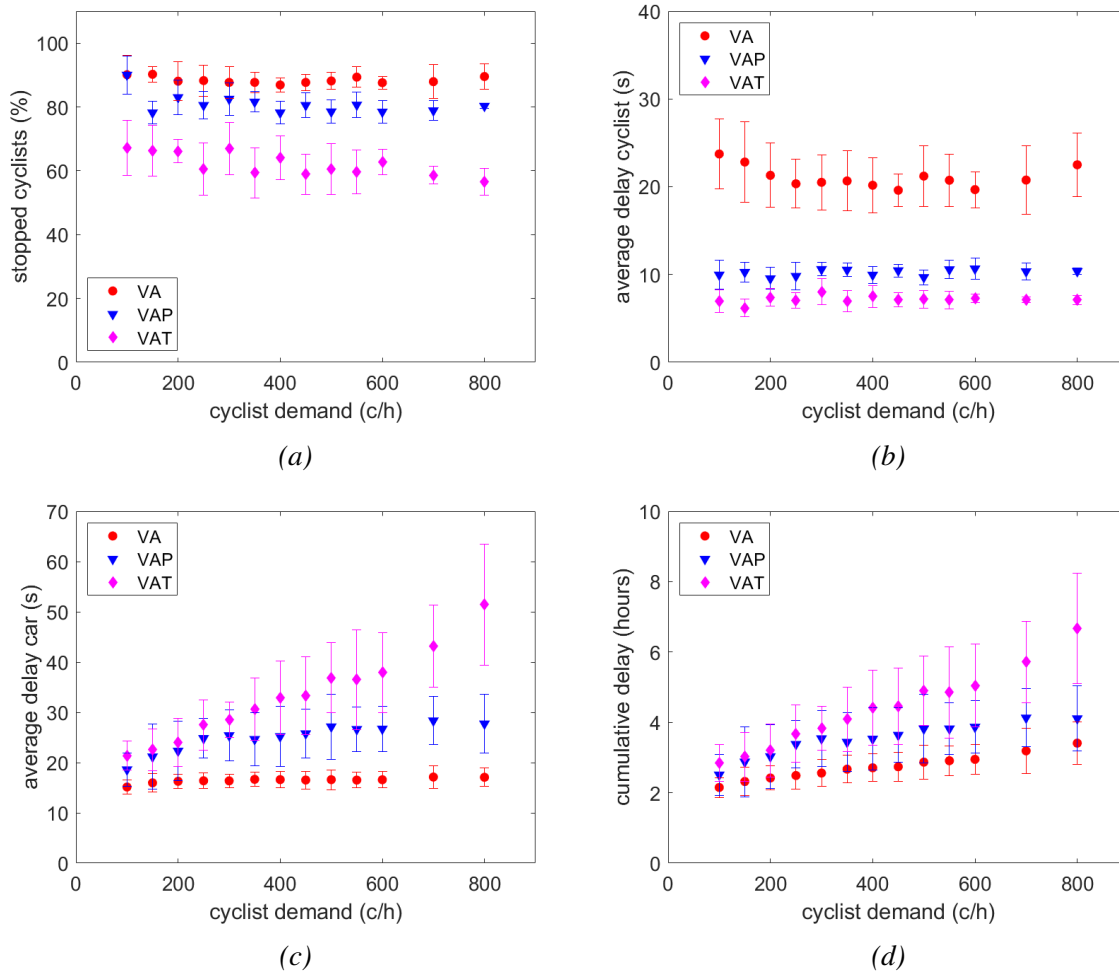


Figure 4.5: As function of the cyclist demand: (a) stopped cyclists, average delay of (b) cyclists (c) cars, (d) cumulative delay.

When the average delay for cyclists is compared to that of car delay, see Figure 4.5c, the effect is smallest for a small demand, but differences get bigger with larger demand. From 400 cyclist/hour, it can be observed that the average car delay start to increase significantly, while the average delay for cyclists hardly changes for higher demand as shown in Figure 4.5a. With increasing demand, the priority for cyclist is given more often and for a longer duration, thus leading to higher delays for cars.

The car demand for the results in Figure 4.5 is fixed. Also the effect of the demand of cars (hence the number of CAVs) is tested for the VAT scenario. The ratio between the demand of car movements remains the same as for the original simulations, but the demand is set to a percentage of this baseline demand (100% equals the baseline demand). As can be seen from figure 4.6a the delay of the cars is strongly related to the car demand, with increasing demand the delay increases much, especially if there is also a high bicycle demand (from 60 to 100 s). For the cyclists this effect is less, because the cyclists receive priority, even when there

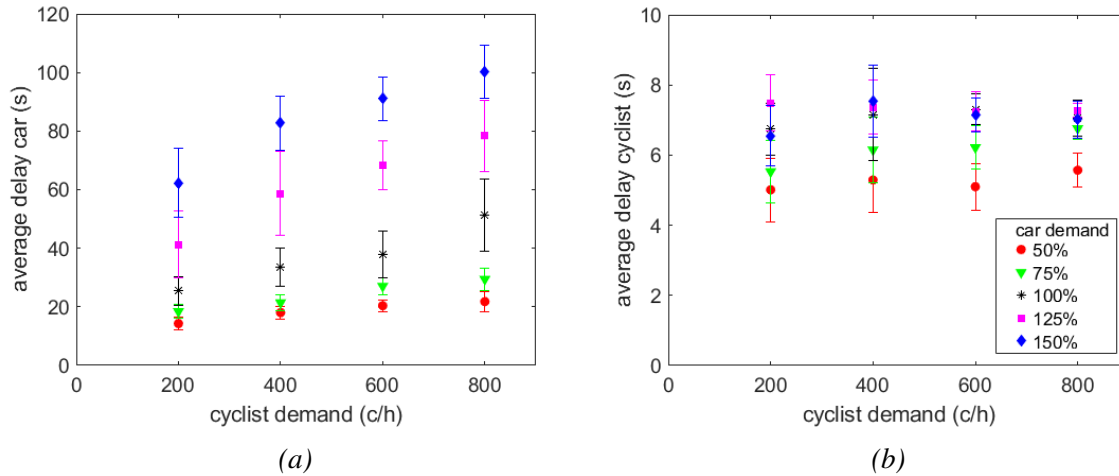


Figure 4.6: The effect of the car demand, as percentage of the baseline car demand (a) average car delay and (b) bicycle delay demand.

are fewer CAVs around to observe them. Once a cyclist is observed it will, in most cases, be “remembered” by the tracking algorithm and will receive priority. Furthermore, if there are fewer cars, the green time of the car movements will be less, cyclist will receive green earlier and the delay will be lower. Only for a demand of 200 c/h the effect of a higher car demand (more CAVs) is visible: with 150% the delay is lower than for 125%.

#### 4.7.2 Q2: bicycle green actuation, bicycle and car green duration and green use

This section aims to answer **Q2** by showing how VA, VAP and VAT each handle bicycle and car green, and the result it has on effective green use for cyclist and car delay. The green handling in Figure 4.7 shows (as function of the bicycle demand) how many cyclists move over the stop line during green, the average green duration for the bicycle signal, and how much time was given per cyclist, that is, per actuation the green duration divided by the number of cyclists during that move over the stop line. As can be seen from Figure 4.7c the green allotted per cyclist reduces as the cyclist demand increases, so the green time is used more effectively with increasing demand, until about 500 c/h where it starts to level. For 300 c/h and more, the green time per cyclist is much higher for VAT control than for VA and VAP, so the green time per cyclists is used less efficiently.

As can be seen in Figure 4.7, VA control does not actuate the bicycle signal often (4.7a), and for a short duration (4.7b). With increasing demand, the green distribution over the cyclists (is the amount of green per cyclists) decreases (4.7c), so for high demand many cyclists pass the stop line during green.

VAP actuates often, therefore it needs lower green times than for VA, since smaller queues can be discharged faster. Since the duration of a signal is at least the minimum green time, this can lead to unused green after the cyclists have moved over the stop line, so the green used per cyclists is higher for VAP than for VA.

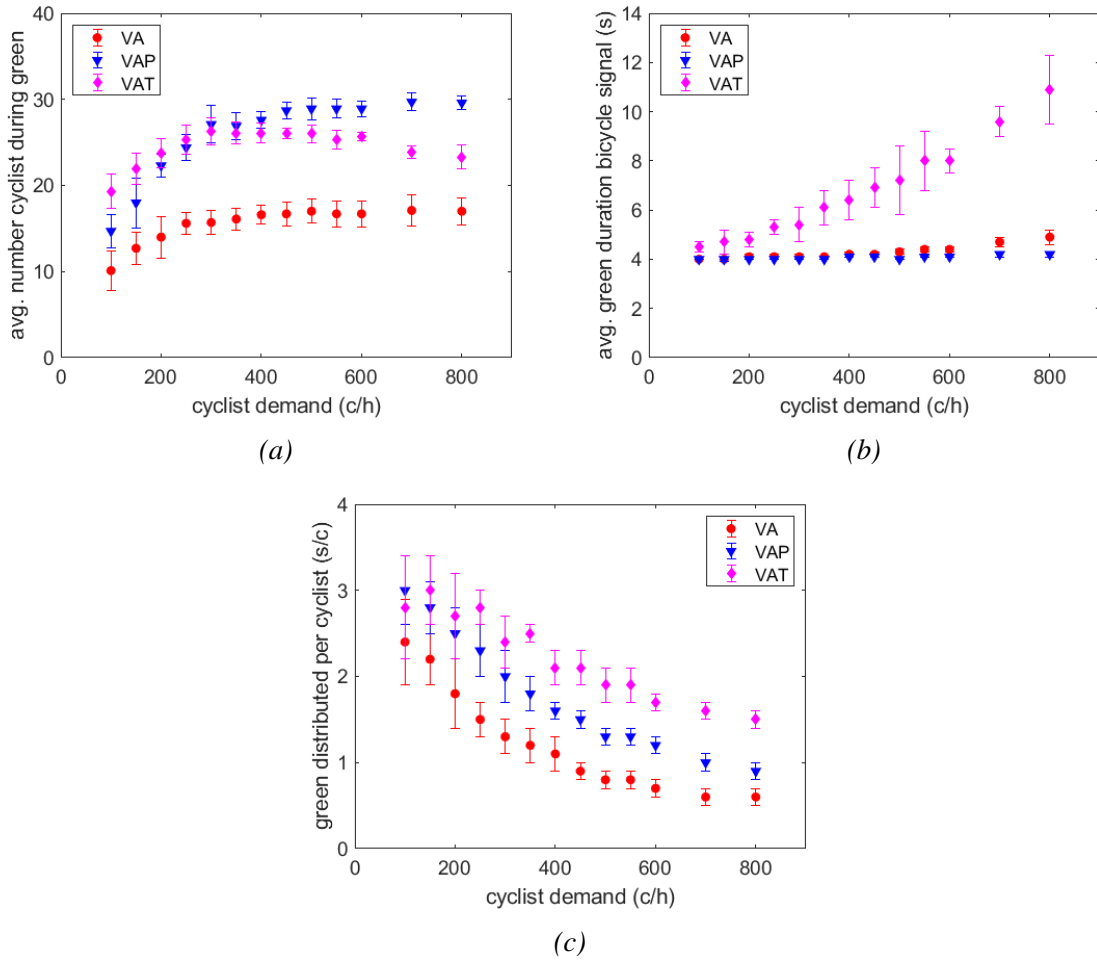


Figure 4.7: Bicycle green handling for VA, VAP and VAT: (a) average number of cyclists that move over the stop line during green, (b) average green duration of the bicycle movement, (c) average green time allotted per cyclist.




Figure 4.8: (a) Average car delay as function of the average bicycle green duration (b) comparison between the desired green duration and the realised green duration of car movements, as function of the percentage of car demand, for different levels of bicycle demand (200, 400, 600, 800 c/h).



VAT also actuates more often than VA, less than VAP, but extends the green time for longer periods, especially when the demand increases, the green is extended to cater for approaching cyclists. Unused green time is present before cyclists move over the stop line. In further development of the controller, this extension of the green should be assessed, since for VAT this results in high car delays (4.8a). This relation is absent for VA and VAP.

In Figure 4.8a the average car delay is given for the average bicycle green duration, for the demand. Results indicate a linear relation between the green time of the bicycle signal and the car delay. In 4.8b, for the different percentages of the car demand, the green time needed to cater for this demand is given. This green time is calculated, by using the Webster cycle time (Webster, 1958), which is optimum for non-uniform arrivals. The Webster cycle time takes into account the control structure and the car demand, and the green time used for the bicycle movement. As can be seen, for equal or more than the base scenario demand, the desired green time is higher than the actual green time the car movements receive, leading to a higher delay, as has been shown in Figure 4.8a.

### 4.7.3 Q3: impact of CAV penetration rate and field of view

The results presented in Figure 4.9 show the difference of weighted average delay between the new scenario and the base scenario. The average is weighted for the number of vehicles, which are different for each simulation. When the difference is positive, the new scenario will give a higher delay than the base scenario.

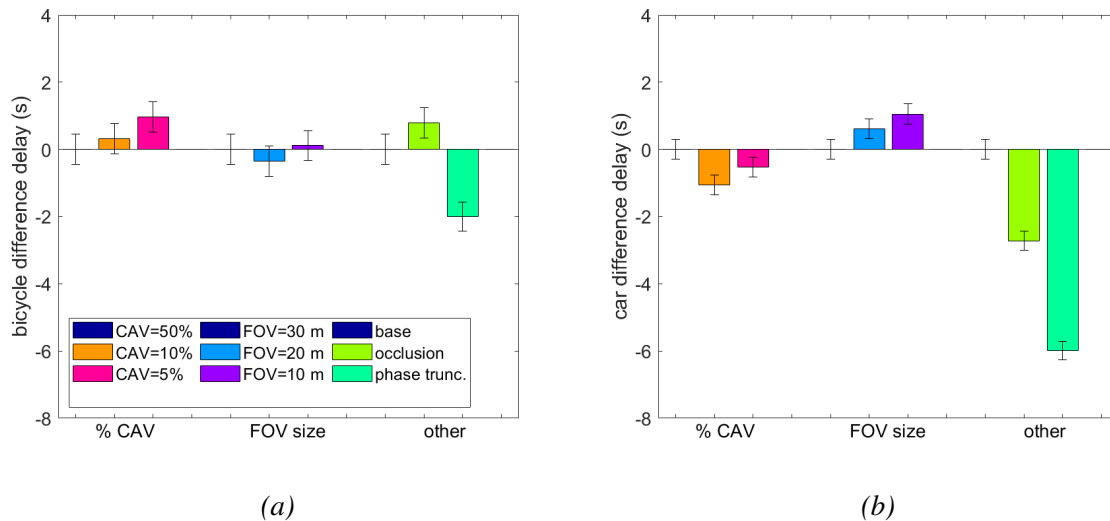


Figure 4.9: Difference with respect to the base scenario for (a) delays for cyclists, (b) delays for cars.

In answering Q3, Figure 4.9 indicates that the penetration rate reduction to 10% has hardly any effect on the delay, the standard deviation is larger than the difference. For a CAV penetration rate of 5% the delay is higher for cyclist. Thereby, the FoV appears to have hardly any difference, a smaller FoV will yield to a bit more delay on average, but the standard deviation is larger than the difference. Note that the FoV chosen for this study are believed to be conservative values as compared to available setups in practice, which benefit from much wider

FoVs. Overall, and as expected, both small FoV or small penetration of CAV leads to less cyclist observations. Because a cyclist observation initiates the creation of a hypothesized track, when not assigned to an existing track, that cyclist's trajectory will be estimated until the stop line to eventually request for priority. The un-interruption of a track may, however, depend on the configured track pruning strategies applied with the tracker. Therefore, if more observations are generated, by a larger penetration rate or larger FoV, cyclists will benefit from the priority requested by a cyclist that has been observed. If it happens less often because of a low penetration rate of 5%, the priority is given less often and the delay will increase with respect to the base scenario.

#### 4.7.4 Q4: what is the value of observing the absence of cyclists?

In answering **Q4**, two scenarios were computed that are representative for situations where cyclists are not observed. Figure 4.9 shows that in the case of occlusion, the delay is higher on average. The tracker removes existing tracks when too much in doubt, thus leaving the cyclist as “unseen” and having no base to request priority. A more detailed look at the behaviour of the tracker confirms that track removal is a consequence of occlusion events in the simulated traffic situation. Thereby, when a CAV generates no observation about a cyclist (because of occlusion) that should have been seen, the likelihood of that (unobserved) track decreases. While the tracker can appropriately handle a reasonable amount of occlusion given the provided occlusion model, occlusion levels increase significantly close to the intersection. A solution to this would be to implement a more sophisticated place-dependent occlusion model, which handles various levels of occlusion at different locations. With more available observations, however, even higher occlusion levels are handled well by the used tracker, as indicated in Figure 4.4.

The second scenario presented here is the truncation of green phase. When the minimum green phase, yellow phase or clearance time is truncated, the gains will be highest. The cycle will be shorter without unused lost time, and priority will be granted earlier, leading to less delay and fewer stops. As can be seen in Figure 4.9, except for the phase truncation, the difference in car delay is opposite to that of cyclist. Where cyclists have more delay, the cars have less, although for most cases the difference is smaller than the standard deviation. Both cars and bicycles gain from phase truncation, smaller cycles with less lost time is beneficial for both traffic modes. Overall, although phase truncation leads to a low delay for all traffic modes, this can only be done when it is absolutely certain that no cyclists are observed within the dilemma zone, and no cyclists are observed in the conflict area, since the risk is very high for a cyclist that is present yet has not been observed. The prevention of occlusion events and false observations, *i.e.*, through redundancy of measurements, is thus an important safety factor.

## 4.8 Conclusion

This work highlighted the potential of cyclist observations made from moving sensor platforms to augment the situational awareness of traffic signal controllers. A simulation study quantified the benefits of the proposed approach, using a vehicle actuated controller that uses CAV observations to prioritize cyclists and reduce lost times in the control cycle, with two baseline scenarios; with and without cyclist priority.

In terms of car and bicycle delay and stops, results indicate that with a low penetration rate and with occlusion, prioritizing cyclists by tracking (with a small FoV) can reduce the delay for cyclists or avoid stops, using a simple algorithm. However, the average car delay increases much when the number of cyclists increases, while the average bicycle delay does not show much variation. Prioritization schemes applied in real-world traffic situations confirm this result, *e.g.*, in the city of Rotterdam (NL), when cyclists were prioritized during rainfall, the queues for car directions increased substantially (Rubio, 2015).

An analysis of efficient use of green for cyclists indicates that unused green for the bicycle movement, or many priority request by cyclists, lead to short green times for cars. An explanation for this result is that the controller does not optimize the delay for all traffic participants, and neither takes into account the platoon dispersion of cyclists.

Both small FoV or small penetration of CAV leads to less cyclist observation. Results highlight the importance of handling false observations as they can lead to unjustified green starts where no cyclists are present. Highly occluded areas (*i.e.*, at the intersection) can have negative impact on the information available to the tracker and thus lead to, *e.g.*, pruning tracks of existent cyclists, which eventually leads to less priority requests.

Results indicate that observational confirmation that cyclists are not present is beneficial for all traffic, leading to less lost time in the cycle.

## 4.8.1 Future work

Future work of this study can span out in multiple directions. We outline a set of interesting paths for future research and discuss practical difficulties implementing such a method.

Although results from the simulation indicate that also a low penetration rate of cyclists can benefit from priority by CAV observations, in practice, things might look a bit different. For instance, we expect control cycles where cyclists can be prioritized, and control cycles where, due to absence of CAVs (in cases of current low CAV penetration rates (Statista, 2022)), cyclists are not prioritized. This may raise false expectations about such a prioritization scheme among cyclists which, in turn, may affect the credibility of the control and may eventually induce red light running (Meel, 2013).

Furthermore, the level of false observations and the avoidance missed detections, *e.g.*, due to potential occlusion events frequent in urban areas, should be tackled as this may lead to losses and delays that are not necessary. The model used to realize occlusion is a very simple model to apply. Further research in literature and by measurements is needed to determine how severe occlusion is when CAVs observe cyclists, depending on location and time, *e.g.* place-dependent occlusion. Further, the value of information in time or quality, *e.g.*, in what way the CAV communicates the quality of its observation in case of occlusion, needs to be investigated.

Experiments under real-world conditions must be performed. Generating real data will provide valuable insight about how efficient CAVs can be in observing cyclists, and enable the calibration of a microsimulation. Furthermore, the communication between iTLC and CAVs must be further developed, extending existing setups where a few static sensors provide traffic data to iTLC.

Cyclists prioritization should be assessed with respect to safety. When a saturated car queue moves over the stop line, ending green can be unexpected by the car drivers, and less credible control can lead to red light running.

The proposed truncation of phases, especially the clearance time truncation, is not fit for use in practice yet. The certainty no cyclists are on the conflict area should be very high before the clearance time can be ended, so first this part of equipment and model should be tested thoroughly before it is used in practice.

The method of prioritization used in current controller, can lead to a situation where one cyclist can get one second less delay and does not have to stop, but this can make, say, three cyclists just arriving at start red of the bicycle signal to have to wait, say, nine seconds (the time for conflicting movements to start and clear the conflict area). This leads to a large variation from simulation to simulation in number of stops and delay. Now that the proof of concept is given that cyclists benefit of prioritization by observation, the next step is to optimize the control, that balances the delay and stops of cyclists and cars. Suitable for vehicle-actuated control is a function that determines the signal phases with an objective function, based on a weighted combination of the total time spent (TTS) by the motorized vehicles and by the bicycles in the system, and their stops.

When cyclists are prioritized, it is beneficial for the throughput of the intersection that the cyclists are close together, so the green extension can be shorter, as is shown at the municipality of Hengelo (NL) where cyclists are encouraged to form platoons so that they can be prioritized (Crow, 2022). Furthermore, a combination of speed advice for cyclists (Dabiri & Hegyi, 2018) and prioritization should also be researched in future.

# Chapter 5

## Privacy challenges introduced by intelligent autonomous systems

---

Chapters 2–4 explored the system’s functional requirements, data integration methods, and traffic-related benefits, while briefly acknowledging privacy concerns without examining potential limitations or harms inherent to the IAS infrastructure. This chapter suggests pressing privacy implications of anticipated changes brought by IAS, highlighting the need for interdisciplinary research and collaboration.

This chapter builds upon the conceptual framework presented in Chapter 1.5, which provides a solid foundation for clarifying terminology and understanding interactions between system components. It shows how the relationship between the machine and the environment is different for IAS than from other existing systems, and how platforms deployed on top of the infrastructure eventually lead to an increased control over the environment to enable the safe functioning. The contributions of this chapter pave the way for further research into comprehensive privacy analyses, including the evaluation of existing privacy frameworks and models in the context of IAS.

This preliminary research provides strong signals to public and private stakeholders by putting on the map a set of novel potential limitations and risks led by the deployment of these infrastructures in environment inhabited by people.

This chapter represents unpublished and preliminary research informed by earlier discussions between Vial A., Gürses S., Troncoso C., and Hoogendoorn S.P.

---

## 5.1 Introduction

Unmanned intelligent autonomous systems (IAS) are systems that are designed to be capable of operating without human intervention. Visions for unmanned IAS in urban environments include self-driving vehicles, drones and other types of connected robots. These different types of systems have in common that they need to continuously sense the environment to function, and to do so safely. This means that IAS, by design, include many sensing facilities. Justified by safety requirements (Sanchez et al., 2016), they are expected to coordinate data across time and space, introducing a mobile surveillance infrastructure into the environments that they function in. The mobile surveillance infrastructure consists of mobile nodes connected through a complex network of computational entities, continuously sharing sensor data with nearby edge-cloud servers. These also include platforms that are essential for the provision of economically viable and scalable digital services. These platforms are envisioned to be used for "measuring, targeting, addressing, formatting, calculating, standardizing, projecting and forecasting" public environments (Parikka, 2023). We foresee this newly introduced infrastructure to reconfigure the lived environment and see these developments as signaling an infrastructural build-out, the consequences of which will have privacy implications. Some of these privacy implications will inevitably enter the field of traffic engineering; changes in traffic flows will need adaptation due to introduction of IAS, while the field will be able to innovate to the science and practice utilizing the sensing data collected for the functioning of these systems. The main objective of this chapter is to stimulate further research into the privacy implications inherent to the infrastructural shift<sup>1</sup> introduced by IAS that arise in traffic engineering use cases, with a specific focus on the tracking of people that inhabit these environments.

We focus on connected autonomous vehicles (CAVs), a specific instance of an IAS with sensing requirements<sup>2</sup>. With the development of CAVs, a variety of applications with significant implications on policy, society and the urban environment, are expected to be deployed (Milakis et al., 2017, 2018; Soteropoulos et al., 2019)<sup>3</sup>. These applications promise to bring benefits to users, including profound changes to people's mobility (*e.g.* self-driving taxis), access to logistics (*e.g.* delivery robots) or data collection capabilities (*e.g.* traffic or environmental

---

<sup>1</sup>The infrastructural shift has already started with the use of cloud technologies, *e.g.*, Amazon Ring's surveillance network and the platformization of services (Stahie, 2019) is expressed by the rise of a digital infrastructure, which can be deployed onto conventional infrastructure (*e.g.*, road or transport networks) with the promise of making these physical structures *programmable*. These *computational infrastructures* consist of a vast global network of data centres, network infrastructure and connected devices, as well as software platforms. The promise of programmability is that by adding a digital layer, the plans and policies of conventional infrastructures can be abstracted from their underlying physical constraints, that is, can be updated, automated and optimized just like digital systems. As part of this process, these environments are instrumented in order for digital devices to function autonomously and safely within them (Toh et al., 2020).

<sup>2</sup>First examples of the rapid development of surveillance infrastructure using CAVs include individual tracking (Greenberg, 2019), mobile cameras for police departments (Gordon, 2022), or the surveillance of distancing between people using camera equipped cars (Bezemer, 2020). In a recent article, the Electronic Frontier Foundation has documented the re-purposing of autonomous vehicle infrastructures for investigation purposes to serve law enforcement requests (Guariglia, 2023)

<sup>3</sup>We highlight the work of (de Sio et al., 2023), which relates values and actions of AVs; sensing is one of those actions, for which the paper posits criteria of tracking and tracing for the AV to satisfy principles of meaningful human control.

sensing). By *users* we mean people who use a CAV operator's product, platform or service. Note that a user can occupy different roles, for instance, a person transported in a self-driving car, an intelligent traffic signal controller receiving information about observed cyclist traffic, or a traffic engineer collecting traffic flow data in a sensed area.

Traffic engineering is crucial to the successful introduction of CAVs onto roads (Tettamanti et al., 2016; Mladenovic & McPherson, 2016; Toh et al., 2020), and the field's transformation with the availability of granular sensing data is likely to further entrench privacy issues as something the field will have to grapple with. The rise of privacy issues in traffic engineering have been longer in the making. In the past two decades traffic research and practice have witnessed an increased availability of various types of traffic data<sup>4</sup>. The large deployment of sensor networks in urban environments and the rapid advancements of mobile and communication technologies have, in another infrastructural shift, enabled the use of different sensing technologies for the collection of information about pedestrians and cyclists. Using the data from dedicated or personal sensing systems has allowed traffic research to empirically study human movement behaviours and to develop, calibrate and validate micro- and macroscopic models. Further, it has pushed traffic engineering practices to deploy more modern traffic monitoring, control and management applications, *e.g.*, evolving from fixed traffic signal timing settings to intelligent traffic light controllers (van der Vliet et al., 2016) using information from approaching traffic via on-board computers, mobile apps (Verbeeke, 2020) or enabled by Radar on a fixed location (Lai, 2021).

In particular, in this chapter we focus on location data, as CAVs are likely to collect and share location data impacting users of these services and this data is likely to trigger the interest of traffic engineers. Efforts of location privacy scholars have been instrumental in understanding inherent risks of localization and (re-)identification from users traces. This line of research gained popularity with advances in another infrastructural shift; the roll out of cell networks and broad use of Wi-Fi enabled the use of a wide range of location-based services. (Krumm, 2007; de Montjoye et al., 2013; Gambs et al., 2014; Zang & Bolot, 2011; Pyrgelis et al., 2017; Xu et al., 2017), propose different location privacy preserving mechanisms (LPPM) to mitigate these problems (Krumm, 2009; Primault et al., 2019; Andrés et al., 2013; Takagi et al., 2020), and introduce metrics to evaluate the proposed approaches (Wagner & Eckhoff, 2018; Zhao & Wagner, 2019). However, their application in the currently developed infrastructural shift enabled by IAS, and here considered with the introduction of CAV, is not straightforward.

With the advent of self-driving cars, drones or other types of connected robots, the amount of sensed features in public spaces is expected to drastically surge. We speculate these types of IAS to become a new source of data for science and practice, in particular, for design, planning, decision-making and innovative services. This raises broader issues of how traffic engineering might be used to optimize physical environments for the seamless and safe functioning of these autonomous systems (*e.g.*, CAVs)<sup>5</sup>.

---

<sup>4</sup>Legislation and regulation have also increased, limiting the storage of privacy sensitive data.

<sup>5</sup>For example, by pushing out pedestrian access to remove complexity and make overall behaviour of traffic more predictable, or requiring a micro-transaction to insure safe passage when crossing the road (Bogost, 2016).

<sup>6</sup> It also raises a plethora of privacy concerns for people who inhabit the environment captured by this infrastructure, which have not yet been studied in great detail by privacy and traffic engineering communities alike.

“Autonomy” and “intelligence” are the two most important features for the seamless and safe functioning of unmanned IAS. To act safely in the world, CAV aim to capture relevant objects in the environment, which inevitably includes people who are not users of the IAS service (non-users). CAV researchers and designers argue (Sanchez et al., 2016) that this is a necessary precondition, and sensing surrounding humans a primary function, to maximize the safety of all traffic participants. In particular, CAVs must observe, model and anticipate human behaviour (Ohn-Bar & Trivedi, 2016) inside and, more importantly, outside the vehicle. Perception sensors provide highly accurate measurements in a continuous stream of data. The fused sensor data is integrated via distributed computational infrastructures in real-time into an environmental model of a CAV’s surrounding, *i.e.*, a digital representation of the sensed environment including roads, objects and free space. Captured objects can be further classified to distinguish between, *e.g.*, other vehicles, pedestrians, cyclists or solid obstacles, producing data that is relevant for — the vehicle’s own functioning aside — among others, traffic engineering.

Since the distinction is important from a privacy perspective, we use the term *non-user* to refer to someone who is in the environment of a CAV and captured by its sensing devices. We assume these are not active or registered users of the service that captures their presence and behaviour. We characterize a non-user as someone that is not offered mechanisms to be informed about the collection of data about them from an operator’s platform or service, and similarly with no options to control nor to protect that data by design, which includes data minimization and privacy respecting defaults.<sup>7</sup>

The privacy of non-users (*e.g.*, pedestrians, cyclists outside the vehicle) in CAV settings should be a major source of concern, yet is a blind spot in the related literature. Existing approaches in CAV and location privacy literature fail to describe the role of non-users, beyond their indication as moving objects in the scene, and hence miss to study the pressing location privacy concerns of this extended sensing setting. Efforts on vehicular location privacy focus on users, and are mostly directed towards communication networks and protocols, proposing schemes to protect and evaluate the privacy of broadcasted messages containing sensitive information (Eckhoff & Sommer, 2014; Zhao & Wagner, 2019; Asuquo et al., 2018; Nguyen et al., 2022). The used models, further, rely on discrete events and leave continuous capture of the environment and the possible tracking of non-users unaccounted for.

Fortunately non-users in a CAV context have attracted some interest in qualitative studies (Bloom et al., 2017) focusing on the expectations and concerns of people surrounding CAVs towards the technology. The capture of the environment by sensing devices including so-called bystanders can be found in literature from other fields, *e.g.*, in the context of life-logging devices, augmented reality technology or camera-based assistive devices (Hoyle et al., 2014; Den-

---

<sup>6</sup>As (Forlano, 2019) notes, safety goes hand in hand with economic development and for bolstering the automotive industry, while narratives of technological determinism around safety testing has shifted, as the many stakeholders advocating for autonomous vehicles argue that the technologies must be tested in the real world to prove their ability to function safely (Leonardi, 2010).

<sup>7</sup>One could argue that non-active users of the service, whose presence can be captured by a CAV, may be provided some of these controls, but this does not change the problem with respect to the limitations of LPPM and also the need to be tracked by all IAS services to be provided privacy protections.



ning et al., 2014; Ahmed et al., 2018; Akter et al., 2020). These works, however, focus on responsiveness, acceptance and behaviours towards these devices, while they entirely ignore location privacy challenges of non-users.

We expect the capture of non-users location data through types of IAS such as CAVs to become highly precious for private companies, policy makers, and researchers, which as described in (Calacci et al., 2019), are already being widely exploited both in a benign and surreptitious way in many of today's mobile applications. The ability to sell and purchase location data or other auxiliary information has become a market practice (Cox, 2022) and frequently the topic of concern in collaborations between tech companies and law enforcement (Guariglia, 2023).

In the following, we develop a cohesive description of IAS and other established digital systems, grounding our analysis in the conceptual framework presented in Chapter 1.5 (Jackson & Zave, 1995). We characterize the data typical for these settings, and make an attempt to formalize their utility in order to facilitate qualitative comparisons across the systems at hand. We then speculate on IAS, using CAVs as concrete use case, *i.e.*, pedestrians and cyclists captured by CAVs to function autonomously, and analyse its components, mechanisms, and interrelations to help distinguish it from existing systems and position IAS appropriately in the field. Building on the developed understanding of IAS and its place within the existing digital landscape, this chapter then presents future research directions, pushing towards more comprehensive privacy analyses. This chapter makes the following contributions:

- Show how the relationship between the machine and the environment is different for IAS than from more traditional Location-Based and Mobile Crowdsourcing Systems
- Identify potential location privacy implications for non-users that stem from the infrastructural shift introduced by IAS entering physical space

This work analyses privacy for non-users and proposes a solid foundation on the issue, for both the privacy and transportation community.

## 5.2 Comparative analysis with established systems

Over the past decade, researchers in the privacy community have developed mechanisms and metrics aiming at mitigating and evaluating location data-related privacy risks. However, existing frameworks and models are designed with Location-Based Systems (LBS) and Mobile Crowdsourcing Systems (MCS) in mind. Unlike IAS, their function does not depend on sensing infrastructure capturing people that inhabit the environment of the mobile device while not being their users. To better measure the effectiveness of existing privacy frameworks in the future, it is essential to understand the key differences between IAS and established digital systems like LBS and MCS. Building on the vocabulary and concepts of the framework presented in Chapter 1.5, the following provides a comparative analysis of how data is represented and utility is derived in each system.

### 5.2.1 Location-based systems

Mobile or vehicular LBS use location awareness to provide users of mobile devices with specific services and location data mining algorithms to determine points of interest and traffic patterns (Andrés et al., 2013).


Typically, LBS are initiated via user location requests, *e.g.*, to find nearby points-of-interest (PoI), check-in restaurants, emergency call, or informing about an accident. A user's location is reported leveraging the infrastructure's positioning system, *e.g.*, GPS satellites, cellular base stations, or Wi-Fi routers distributed on mobile devices, *e.g.*, smartphones, wearables, or vehicles. The interaction with the environment is limited to input from and to the user (based on the location of the user). The machine changes state based on inputs from user and the location of the user, and eventually effect the output. The generated location data enables a service provider to track users across location, giving the service provider the possibility to build services based on locations of many users (*e.g.*, traffic monitoring), which may also effect the environment (*e.g.*, users may go to certain PoI, or change paths to avoid traffic).

**Data** – The state of a user at a certain event is typically modelled as a triplet (Shokri et al., 2010),

$$X^{\text{LBS}} = \langle \text{identity}, \text{position}, \text{time} \rangle.$$

The *identity* of a user can be defined by any feature that makes a user distinguishable from another. This distinguishability provides the distinct association between a user and the *location* information of a request in *time*. Typical LBS use single locations, *i.e.*, revealing the current location once to a service provider to obtain a service (*e.g.*, “nearby PoI”), whereas other LBS report locations more often resulting in trajectories (*e.g.*, routing and navigation services, or sports tracking). The time distribution of reported information can thus vary between snapshot, sporadic, or continuous (Shokri et al., 2010; Asuquo et al., 2018), where the latter is more common in vehicular LBS.

**Utility** – Let  $U = \{u_n\}_{n=1:N}$  be a set of users, where  $N$  is the number of users. At time  $k$  the state  $X_k^{\text{LBS}}$  of a user  $u_i$  is shared upon request of the user's mobile device with service provider  $SP$ . In return, a utility is obtained by the user at time  $k + 1$ . Let  $F = \{f_m^{\text{LBS}}\}_{m=1:M}$  be the set of utilities, where  $M$  is the number of utilities a user can obtain<sup>8</sup>. More generally, we define the user-centric utility to depend on the service query of a user, the granularity and accuracy of the shared location, and can be represented as



Note that, in existent location privacy literature, LBS are typically modelled such that the service provider is not linked to any utility. We argue, however, that the service providers may receive some utility  $g$  in return of the offered service as represented by the dashed arrow. This, for instance, can be in form massive collection of request data which, most often, is exploited

<sup>8</sup>In some cases several utilities can be obtained from sharing location information, *e.g.*, maps, routing, traffic updates, social networking, or context advertising.

for operational improvements, developments of further services (*e.g.*, how busy is a PoI?), and monetization.

## 5.2.2 Crowdsourcing systems

Crowdsourcing systems leverage a platform's available sensing and/or processing capabilities, frequently from an unknown number of participants, to solve tasks in a distributed manner. Mobile and (manned) vehicular crowdsourcing applications are mostly location aware (Chatzimilioudis et al., 2012), and capture single or specific indicators collectively using positioning, sensing and processing infrastructure. In contrast to LBS, the interaction with the environment is based on user interaction. These systems typically link the location data with measurements of application-dependent sensors. The machine thus changes state based on measurement and location of each user. The location and sensing data enable provider to build services that aggregate location and measurements of many users (*e.g.*, pollution sensing). In other words, MCS execute data collection service via the platform — we like to call, observations as a service (OaaS). The results may impact the environment, *e.g.*, user behaviour, how you organise resources to clean up the pollution. The geo-located measurements are typically shared with servers to attain a particular global objective that can, in some cases, generate the foundation for LBS applications. For example, Google Maps relies on crowdsourcing for the generation and updating of their maps, while these maps can then be used to provide location-based services to users (*e.g.*, routing).

**Data** – Similar to LBS, crowdsourcing applications typically include *location* and *time* information. Crowdsourcing settings differ from LBS as they may report a *measurement* from various types of sensors, *e.g.*, radiation level tied to the timestamped location of a user (Safecast, 2022), or acceleration measurement from a smartphone (Matarazzo et al., 2018). The data can be modelled as

$$X^{\text{MCS}} = \langle \text{identity}, \text{measurement}, \text{position}, \text{time} \rangle,$$

where the type of sensor used to collect measurements depends on the utility of the application. In contrast to LBS settings, however, data is typically collected continuously over longer periods of time. Note that MCS leverage the positioning infrastructure, also used in LBS, to link measurements with location data and turn them into crowdsourced information. For example, accelerometer data collected from smartphones can be linked with positions and provide valuable information on traffic conditions (Mohan et al., 2008).

Collected data points can be linked via personal user identifiers often present in more traditionally LBS settings, or, more indirectly, using a quasi-identifier obtained by combining each user's (unique) device identifier, the location and other available information (Boukoros et al., 2019).

**Utility** – In contrast to LBS, and as described in (Boukoros et al., 2019), the utility cannot be captured with a user-centric approach as crowdsourcing benefits from aggregating data collected by a large amount of users. Let thus  $Z_k^V$  be the set of all  $M_k^V$  measurements received by the MCS operator from a subset of data collectors  $V$ , where  $V \subseteq U$ , such that  $Z_k^V = \{Z_k^{u_1}, \dots, Z_k^{u_{M_k}}\}$ , and where  $Z_k^{u_1}$  includes  $X_k^{\text{MCS}}$ , *i.e.*, data received at time  $k$  from user  $u_1$ . Further, let  $W$  be a set of users that are not data collectors, where  $W \subseteq U$  and  $W \not\subseteq V$ , who can obtain utility without collecting and sharing measurements.

$$\begin{array}{c}
 \xrightarrow{f_{k+1}^{\text{MCS}}} \\
 U \xrightarrow{z_k^V} SP \\
 \xleftarrow{f_{k+1,j}^{\text{LBS}}}
 \end{array}$$

This form of collective and value-based utility  $f_{k+1}^{\text{MCS}}$  exists through users (here represented as dynamic set  $V_k$ ) that act as data collectors, however, applications may append LBS-like user-centric utilities  $f_{k+1}^{\text{LBS}}$ , once enough data is available (*e.g.*, openstreetmap). In other words, the aggregated utility for the users can be the same as for the service provider (and depends on the number of participants?).

### 5.2.3 Intelligent autonomous systems

IAS include, based on notions introduced in Chapter 1.5, all the users, the respective surrounding environment, as well as the devices (*e.g.*, CAVs) and server infrastructure of the utility/sensing provider.

Jackson & Zave (1995) avoid to provide a literal definition of the *environment*, but describe it informally as mechanisms of the (mechanical) apparatus and its use (by users), where *environment phenomena* can be understood as predicates characterizing sets of events. We complete this informal definition in three ways: (1), we introduce a dynamic dimension including static and dynamic objects (*e.g.*, roads, animals, and non-users), local conditions (*e.g.*, climate, traffic, and network topology) and the own location of the mobile device as input that may impact its safe functioning. The dynamic nature of environment states may change the state of the machine, which may require *adaptive* actions and reactions of the system to user input; (2), we expand the temporality of phenomena as we regard events to be successive, but also simultaneously possible. To be *autonomous* and *adaptive* the environment must be sensed (and processed) across infrastructure, possibly sharing data about one or multiple events, while potentially at the same time; and (3), we consider an additional component for the view of time, as phenomena are predictable based on past events. In other words, the potential next state in the environment is predicted based on previously captured data to eventually decide upon the next state of the machine.

Now, let us focus on the interaction between the machine and the environment. According to Jackson & Zave (1995), if the machine is to interact with the environment some phenomena must be shared by both, while machine- or environment-controlled events are initiated by either the machine or the environment. For example, if there are no users interacting with a traditional car, no sensing event (*e.g.*, throttle position) will ever occur, regardless of the machine's behaviour. In contrast, our framework regards the possibility of control to be detached by the physical presence of users<sup>9</sup>, as the function of the machine may primarily be dependent on the environment state excluding (or putting second) user input. As we shall see later, when exemplifying IAS using CAVs, the idea that some environment-controlled events might be constrained by environment properties, and thus the machine can exploit these constraints (through programmability) to prevent events from occurring, is surprisingly accurate.

---

<sup>9</sup>Typically, the environment is where the user is.

In order to provide further detail and highlight the changing role of non-users given previous works, the following section provides more detail into each of these aspects by comparison with existing commonly used systems. To do so we make a set of assumptions about various components of the system, *e.g.*, according to the flow of information, and analyse what sets it apart and enables the shift from common infrastructures enabled by LBS or MCS.

### Connected autonomous vehicles

To make our case concrete we focus on CAVs, a specific instance of IAS where the material infrastructure is exploited via platforms to provide services and empower, for instance, transport applications like self-driving cars. Thereby, the surrounding environment is continuously captured by the infrastructure as a necessary precondition to self-operate in the environment. The available sensing and processing infrastructure enable services based on each CAV to theoretically navigate the environment in full autonomy. Furthermore, the connectivity infrastructure enables communication between vehicles (*e.g.*, via V2X) and to the edge cloud server (*e.g.*, via 5/6G) yielding offloaded application-dependent computation services while maintaining low latency and high accuracy. Localization, perception, tracking, and control services, power a large number of applications while in interplay with the offloaded data and map information. This not only allows to improve situational awareness of their environment, *e.g.*, collaborative object localization at intersections as an important task for the safe and efficient autonomous operation, but also enables a variety of secondary applications, *e.g.*, using data for traffic engineering/surveillance applications.

CAV can include LBS and crowdsourcing services too. For example, a user may request a ride sharing service via mobile phone (*i.e.*, mobile LBS), and eventually be driven by a self-driving car (*i.e.*, vehicular LBS) to the desired destination. This — IAS as means of transport — is what is commonly assumed to be the primary mobility applications of urban CAVs. On the other hand, the utility/sensing provider may facilitate other services based on CAVs to empower secondary applications. For example, a municipality (dependent on the infrastructure of utility/sensing provider) may access observations about different mode choices of people or the green canopy across the city, and thus exploit the sensing done by the infrastructure (*e.g.*, fleet of self-driving cars) for transport, urban planning or environmental monitoring tasks (*i.e.*, MCS).

Table 5.1 provides an overview of main distinctions between IAS (*e.g.*, CAV), LBS and MCS, but also with related existing real-world applications that may, at first glance, appear to be very similar. We distinguish between the surveillance infrastructure enabled by IAS (*e.g.*, CAV) and existing ones in six distinct ways: 1) it is not dedicated to surveillance applications, 2) it includes non-users, 3) user input is secondary, 4) it increases non-user sensing without additional physical infrastructure, 5) it enables responsiveness in real-time and 6) is used to make physical infrastructure programmable.

1. While most LBS and MCS are also not primarily dedicated to surveillance, a misleading analogy can be made with surveillance deployments expanding in public spaces (*e.g.*, CCTV) and on private grounds (*e.g.* Amazon Ring (Stahie, 2019)), typically leaving non-users no possibility to self-control opt-in or opt-out choices (Zhang et al., 2021b). CCTVs, however, are dedicated to surveillance whereas CAVs sense to operate safely in

the environment.

2. The necessity of sensing the environment — ideally from many CAVs — includes human non-users and is what sets it apart from existing systems, *e.g.*, MCS infrastructure that is usually aimed at sensing a single or specific indicator. These (non-users) are considered as one of the environmental objects and differ from users as they have not subscribed to the service nor do they want to be part of it. While this problem may exist in some MCS applications (*e.g.*, Google Maps), the capture of non-users in CAVs, however, is precondition for the safe functioning and operation. Here again, a misleading analogy could be made with existing applications such as CCTV, Amazon Ring or Google Street View, where non-users are captured by a camera at the entrance of a shop, in the garden via the neighbour's doorbell, or in stitched virtual-reality photos done by sensing vehicles walking along the road. These, however, do not rely on the sensing of the environment to function.
3. As the input is aggregated to capture the environment across infrastructure, the input of the user becomes secondary. The input from the environment may change the state of the machine which, eventually, may require adaptive CAV actions as well as reactions of the system to user input (*e.g.*, change how you go from A to b).
4. The mobility of the devices and the existing material infrastructure make it possible to cover the entirety of an area with a single device. Adding more devices, on the other hand, can be achieved without adding more material infrastructure, and can help increase the spatial and temporal granularity in an area.
5. IAS (*e.g.*, CAV) can be distinguished from LBS and MCS as it must react in real-time according to the sensed environment, while in some ways resembles CCTV, *e.g.*, in the context of license plate readers. Using input data from the environment, however, CAVs must predict the next possible state of all things in the environment. Based on this, they change state, adapt, act in the environment.
6. The introduction of computational infrastructure onto our common physical infrastructures make it become programmable. For example, the programmability of the infrastructure further increases the possibility of the machine to control the environment to enable its safe(r) functioning (*e.g.*, traffic lights, road engineering, placement of trees). <sup>10</sup>

---

<sup>10</sup>As the data collected about the user-context is similar to LBS or MCS setting, we refer the interested reader to respective literature that study location-based user-centric and collective value-based utilities. For the rest of this chapter, we concentrate our attention to the capture of the environment by CAVs, in particular, to the role of non-users.

Table 5.1: An overview of the different systems and exemplified applications.


|                             | <b>LBS</b>     | <b>Crowdsourcing</b>    | <b>IAS</b>              |                    | <b>CCTV</b>    | <b>Amazon Ring</b> | <b>Google Street View</b> |
|-----------------------------|----------------|-------------------------|-------------------------|--------------------|----------------|--------------------|---------------------------|
| <b>Target</b>               | User           | User                    | User                    | Non-User           | Non-User       | Non-User           | Non-User                  |
| <b>Deployment</b>           | Individual     | Individual or networked | Individual or networked | Networked          | Individual     | Individual         | Individual                |
| <b>Information</b>          | Loc            | Loc and Meas            | Loc and Meas            | Meas at Loc        | Meas at Loc    | Meas at Loc        | Loc                       |
| <b>Utility (on the box)</b> | User centric   | Collective              | User centric Collective | Collective         | Surveillance   | Surveillance       | Mapping                   |
| <b>Sensors</b>              | Mobile         | Mobile                  | Mobile                  | Mobile             | Static         | Static             | Mobile                    |
| <b>Collection Pattern</b>   | Often sporadic | Regular                 | Sporadic or regular     | Continuous         | Continuous     | Regular            | Snapshot                  |
| <b>Spatial Coverage</b>     | User centric   | User centric            | User centric            | Mobile section     | Static section | Static section     | Mobile section            |
| <b>Data Type</b>            | Tabular        | Tabular                 | Tabular                 | Multi-modal Stream | Stream         | Stream             | Image                     |

**Data** – The position of the CAVs are known given precise positioning systems. When perceiving the environment, a set of measurements (*i.e.*, scan) is generated at the specific point in time that includes static and dynamic objects. In this process, the raw data captured from available sensors' fields of view, is processed through a range of detection, localization and classification assignments to enable the CAV to “see” the environment in which it operates. Typically, and in contrast to LBS and MCS, the capture is a continuous and uninterrupted stream of data, required to guarantee the operation in real-time based on complex data. While dynamic objects can also be other CAVs, animals or robots, we focus on human environmental objects, *i.e.*, non-users, to evaluate their privacy. A non-user can only be observed when in the field of view of a sensor, while is defined in *object* classes (*e.g.*, pedestrian, cyclist) according to a certain likelihood, mostly relying on sensor, algorithmic and computational performance. Humans are therefore reduced to objects, which is underlined in the way their *subject* type is erased. A captured non-user *object*, can be described by its *dynamic state* at a specific *time*

$$\langle \text{dynamic state, time, object} \rangle.$$

The dynamic state vector typically includes *position* and *velocity* and can, for example, be augmented with *acceleration*. Non-users generate data but not from their personal device and their *object* attribute is not linked to a permanent and unique identifier that is kept when an object leaves the field of view. We assume the case where there is no personal device that can be sensed to eventually show that even in this setting there are privacy problems.

**Utility** – As described above IAS may include LBS and MCS applications that provide user-centric  $f_{k+1}^{\text{LBS}}$  or collective  $f_{k+1}^{\text{MCS}}$  utilities for users. We argue, however, that the utility function is augmented to have some primary utility from capturing non-users  $f_{k+1}^{\text{IAS}}$  to enable the main function of CAVs. In fact, applications deployed by service providers (*e.g.*, self-driving) rely on the precise perception of non-users to generate utility for users and themselves.



It becomes clear that the user-centric  $f_{k+1}^{\text{LBS}}$  or collective  $f_{k+1}^{\text{MCS}}$  utilities are conditioned on the primary utility  $f_k^{\text{IAS}}$ . Platforms deployed on top of the infrastructure eventually lead to an increased control over the environment to enable the safe functioning. The expansion of the environment as described in Jackson & Zave (1995) to something more dynamic and less user-centred, which is exposed to constant change and transformation, is inherent of IAS and captured by the proposed framework. The environment is conceived as all sorts of things, which collapses differences between objects and humans. The relationship between the machine and the environment is thus different for IAS than from other existing systems. A thought provoking realization, however, is that the service utility could decrease/or become negative, *e.g.*, if non-users make themselves unrecognizable for sensors (Heathman, 2017), or increase with roads closed to non-motorized traffic. We anticipate that the interplay with material and computational infrastructure will generate tussles similar to what has been observed at the beginning of the twenty first century in internet infrastructures (Clark et al., 2002), ultimately affecting the utility provided to users, while potentially also to non-users.



## 5.3 What's next

The comparative analysis provides a qualitative understanding of key differences between IAS and established digital systems, highlighting the need to investigate if existing privacy approaches apply to non-users, and IAS infrastructure in more general.

A comprehensive privacy analysis of IAS must challenge existing privacy frameworks and models. Theoretical and quantitative evaluation of location privacy preserving mechanisms (LPPMs) is needed, drawing inspiration from Boukoros et al. (2019), while potential extensions could involve replacing Geo-Indistinguishability (GeoI) (Andrés et al., 2013) with Geo-Graph-Indistinguishability (GeoGI) (Takagi et al., 2020). A practical application could be derived from traffic engineering scenarios, *e.g.*, where individuals are captured by CAVs for traffic monitoring, control, and management purposes.

**Toy example** – Building on recent studies (Salazar-Miranda et al., 2023; Vial et al., 2023b), we consider a scenario where CAVs enhance situational awareness for traffic monitoring facilities, previously limited to static sensors. In this context, CAVs monitor temporarily deployed<sup>11</sup> urban road infrastructure, such as pop-up pedestrian bridges or cyclist lanes. *Non-user* are captured by one or more CAVs and can be described by its dynamic state at a given time. Note that despite our understanding of CAV data production, a dearth of information remains regarding its utilization, particularly in aspects like real-time dissemination and offline analysis. Building on Liu & Gruteser (2021) proposal for offloading data capture to servers and shared databases, this assumption seems reasonable given common data market practices, such as Tesla's Autopilot 'snapshots' (Harris, 2022a). From a utility perspective, CAV data can help assess the impact of temporary measures, infrastructure usage, or route choices by measuring pedestrian/cyclist traffic characteristics on alternative routes were previously unobserved.

With this in mind, an adversarial scenario will have to be determined, including both adversary capabilities and goals, to examine the application of traditional privacy frameworks. Using real-world data or simulations, it can be evaluated (with appropriate metrics), how LPPMs apply in providing both utility necessary for the functioning of safety-relevant applications of the system and privacy for non-users. We expect, however, that the quantified analyses will reveal the weaknesses of existing privacy frameworks by showing that existing LPPMs impose unacceptable price on utility and that many of them do not even provide good privacy guarantee for non-users.

Future work could also explore the role of consent solutions in visual systems, similar to those used for online photo sharing, which have proven challenging to implement in non-user contexts (Olteanu et al., 2018). Additionally, investigating the effects of non-user privacy on CAVs operating on dedicated roads could provide valuable insights. Moreover, the research community needs to qualitatively examine the impact of this new mobile surveillance infrastructure on society and urban spaces, acknowledging potential conflicts that arise when technology becomes integrated into infrastructure. We also invite consideration of fundamental questions: can there be utility for non-users, and can one only have privacy if part of the infrastructure?

---

<sup>11</sup>These temporary deployments are emergent as safety and comfort of pedestrian and cyclist spaces influences the quality of life in cities, but crowding can limit these gains (*e.g.*, perceived level of comfort on a busy walking or cycling facility during city events).



# Chapter 6

## Conclusions, implications and recommendations

This final chapter presents the main findings and overall conclusions of this thesis, as well as the implications for practice and recommendations for future research.

### 6.1 Main findings

The main objective of this thesis is to investigate the potential utility of pedestrian and cyclist data generated from moving sensors (IAS or MSP)<sup>1</sup> for advancing traffic research and practical applications. Additionally, it examines the broader societal implications associated with the capture of such data within urban settings. To achieve this goal, four main research questions (RQ) were defined and explored in a series of studies. In this section, the main findings of these studies are highlighted by answering the research questions set out in Section 1.4.

#### **What are the necessary technical and functional requirements to a pedestrian and cyclist traffic sensing system that uses moving sensor platforms to collect data? (RQ1)**

Pedestrian and cyclist data can be used across various active mode research tracks, with studies ranging from comprehensive network analysis to detailed link, or cross-section demanding more or less time-sensitive information. Current data collection methods, however, are often restricted to narrow areas of active mode research because of their inherent limits in spatial and temporal resolution, intrusiveness, and ability to handle dynamic network conditions. Literature suggests that integrating sensor mobility, given a certain quantity of sensors, can enhance the performance of wireless sensing networks. The optimal placement of mobile sensors is crucial in achieving comprehensive coverage and accurate data collection. For instance, network-wide microscopic sensing tasks can benefit from densely deployed sensing platforms with high cov-

---

<sup>1</sup>A reminder that while generation of extensive perception data including pedestrians and cyclists is predominantly propelled by the progression towards IAS, as has been exemplified with CAVs, the generation of data does not mandate that sensor platforms must possess autonomy. Both IAS and human-operated MSP, are capable of capturing information about their surrounding environment. Chapter 2, 3 and 4 present findings applicable to both IAS and MSP settings, with an occasional preference for illustrating one over the other for argumentative purposes. Notably, Chapter 5 focuses on IAS; however, it's important to note that many of its conclusions can be transferred to settings involving MSPs.

erage ratios, whereas local macroscopic sensing applications can thrive with fewer, but more focused deployments. A multi-layered hybrid architecture is suggested as it leverages distributed and decentralized communication and computation, enabling sensor nodes to locally communicate with surrounding ones while efficiently scaling with increasing network size and utilizing parallel processing for real-time data analysis. The multitude of scenarios in which real-world data represents a potential asset for solving research questions necessitates a thoughtful mapping of objectives to sensing requirements and network characteristics, tailored to specific contexts and environments. Overall, it is suggested to leverage current technologies and maximize the novelty of data received at any given moment while avoiding data congestion. With the increasing amount of data being generated, processed and potentially stored, edge computing is a necessary element of the processing layer. Advancements in perception and powerful deep learning approaches have made it possible to accurately capture the environment surrounding moving sensor platforms, however, ethical, privacy, and security concerns are yet to be addressed if used for future traffic surveillance applications.

**What specific adaptations are required for advanced state estimation and data association methods to effectively infer trajectories from data collected by both static and moving network-constrained sensors? (RQ2)**

Inferring individual trajectories opens up opportunities for investigating diverse areas of study within active mode research. A review of literature on state estimation and data association methods suggests that the origins of most these approaches predominantly lie in ground or airborne autonomous systems; domains that have significantly contributed to the development of methods. These approaches, however, typically consider spatially restricted and short time horizons (as they are mainly developed for safety-critical applications), and/or are designed for unconstrained motion in two- and three-dimensional space (*e.g.*, airborne ground moving indicators as sensor). The network-constrained multi-target tracking framework proposed in Chapter 3 fills this gap and extends the traditional target tracking representation with a discrete component placing the target on a given road segment in the network. Simulation results particularly highlight network-constraint effects for efficient target predictions over extended periods of time, and in simplifying the associations process, as compared to not utilizing network structure. Furthermore, the beneficial element of so-called negative information, reporting non-presence of targets, is highlighted and provides convincing evidence on the relevance of all types of data. The versatile framework is applicable across many sensor settings where presence and non-presence of targets are available from a sequence of measurements, and provides researchers and practitioners with a means to draw a micro- and macroscopic picture of the traffic flows in a defined network for a variety of use cases.

**How can this information be leveraged to improve traffic control and safety? (RQ3)**

A review of relevant literature and real-world applications indicates a mismatch of cars and other modes (the study in Chapter 4 focuses on cyclists) data shared with iTLC as a means to optimize traffic at intersections, making it difficult to efficiently optimize traffic control and improve safety. Because existing solutions to improve the situational awareness of traffic controllers come with a considerable cost of deployment and maintenance, Chapter 4 investigated how using cyclist observations from mobile sensing platforms could augment the situation awareness of traffic signal controllers, and thus enable prioritizing cyclist and reduce lost times in an effective way. The study put the framework proposed in Chapter 3 to the test and results from a simulated study highlighted the potential of using CAV-generated cyclist observations to pri-

oritize cyclists and reduce lost times as compared to not using this novel type of information. The study indicates that already a low penetration rate of CAVs, and therefore a limited number of available observations, can reduce the delay for cyclists or even avoid stops, using a simple algorithm. Simulated results also confirmed that observing non-presence is beneficial for all traffic. The study proposes a viable approach for conducting practical experiments leveraging current iTLC infrastructure in any sensing setting that transmit target observations indicating their presence or non-presence.

#### **What are the interrelated potentials and conflicts introduced by the new mobile sensing paradigm? (RQ4)**

Previous studies acknowledged privacy-related ethical issues. The study presented in Chapter 5, however, provides a more detailed assessment of limitations, risks, and potential harm associated with an IAS infrastructure. The privacy of non-users (*e.g.*, pedestrians or cyclists outside the vehicle) in IAS-like setting is a blind spot in both traffic and privacy research literature. The analysis and comparison of existing systems (LBS and MCS) vs IAS shows that the privacy implications stem from this newly introduced infrastructure and holds the potential to reconfigure the lived environment. We see these developments as signaling an infrastructural shift. In contrast to LBS and MCS, the utility function in IAS is augmented to have some primary utility from capturing non-users to enable the main function of the IAS. We thus concluded that the user-centric or collective utilities are conditioned on the primary utility. These findings gesture towards that location privacy implications will inevitably impact traffic research and practice<sup>2</sup>, challenging possible benefits induced by the introduction of IAS (*e.g.*, optimized traffic flows or the availability of novel data). The study outlines the research to be performed, particularly presenting the urgent need for a comprehensive assessment of existing privacy frameworks and models. The findings shall motivate traffic and privacy scholars to work together better understand inherent conflicts implied by IAS-type data collection infrastructures, *e.g.*, related to location privacy, and develop new approaches and mechanisms to better balance utility and privacy.

## **6.2 Overall conclusions**

Based on the chapter findings, this section draws overall conclusions to the main research objective, highlighting the technological foundations that enable utility and the societal implications that emerge from the development of this sensing paradigm.

At the technological level, progress primarily driven by fields related to IAS demonstrate by implementation that available technology, as we exemplified with CAVs, effectively supports key aspects of the proposed sensing paradigm today. This suggests that data on pedestrian and cyclist presence, locations, and movements is being generated, while more or less available, enabling the potential for informed research and applications in active mobility at both macro- and microscopic-levels. Thorough review of existing literature throughout the chapters of this thesis, however, reveals a notable absence of both theoretical frameworks and practical implementations for the sensing paradigm, across multiple scientific disciplines, highlighting the

---

<sup>2</sup>We note that privacy implication within traffic engineering have long been in the making, *i.e.*, mainly because of the increased availability of recent data collection methods, pushing practices to deploy more modern traffic surveillance.

need for further investigation how the data will be collected to substantiate its full benefits. The multitude of scenarios in which real-world data represents a potential asset for solving research questions necessitates a thoughtful mapping of objectives to sensing requirements and network characteristics, tailored to specific contexts and environments. For instance, network-wide microscopic sensing requires high quantities of sensor platforms and large coverage ratios for detailed movement tracking, whereas local macroscopic sensing can benefit from limited coverage for acceptable traffic estimates. Notably, this work suggests that even with limited number of sensor platforms, it may be possible to meet specific microscopic application requirements, potentially at a network level. This underscores the potential of adopting data-minimization principles in sensor deployment, optimizing resources while maintaining acceptable levels of privacy and utility. Furthermore, our work highlights the value of incorporating information on target absence, which can provide critical insights even in challenging environments. To advance this work from its current stage, further improvement of system architectures (*e.g.*, towards more decentralised processing schemes), algorithms, and external control systems is crucial. From a pure technological perspective, these enhancements will be essential for transitioning our work towards practical deployment.

The sensing paradigm presented in this thesis is deeply intertwined with broader societal factors, including technological, social, and institutional considerations that collectively shape distinct usage patterns and raise important ethical concerns. The ownership of deployed infrastructure and associated data creates ambiguity, affecting how data will be collected, managed, analysed, and utilized. The realization that industry will play a prime role in determining architectural design choices, which in turn influence scholarly activities within this domain, must be considered. Furthermore, academic conversations surrounding the privacy concerns of pedestrians and cyclists within IAS still seem to often overlook critical nuances. This work highlights the evident gap in existing privacy and transport literature on this specific subject. Our first examination into location privacy implications uncovered an intriguing perspective, wherein pedestrians and cyclists can be understood as not active beneficiaries (*i.e.*, non-users) but rather as entities within a system that do not directly provide them with services or benefits. This suggests that existing location privacy preserving mechanisms may not function as intended as they are designed for users of a system. This intuition warrants further comprehensive investigation to confirm its validity. This work also suggests broader risks extending beyond individual location privacy concerns; optimization-driven systems, designed to exploit behaviour and environment data for value extraction, and risks related to the concentration of power and control through infrastructure ownership. As we move forward, it is important to consider the broader implications of these technologies and ensure that their benefits are equitably distributed, rather than perpetuating existing power dynamics. Bridging scientific disciplines seems essential to reevaluate existing location privacy preserving mechanisms and their design principles, as well as more nuanced discussions about the responsibilities of infrastructure owners.

### 6.3 Implications for practice

Given the interdisciplinary nature of the research, this thesis offers an array of implications for practice. In the following, we elaborate upon three categories of implications for practice: addressing real-world challenges, the socio-technical understanding of stakeholders and actions aimed at enhancing privacy and security.

### Addressing real-world challenges

To achieve effective problem-solving, we recommend a proactive approach that prioritizes analysing utility and limitations from the outset. Practice is encouraged to identifying the many real-world problems and understanding their context, as solely throwing technology at often misunderstood issues has not proven effective in the past. There are many cases where data is lacking or inadequate, creating opportunities for innovative data collection like the one presented in this work. For example, collecting data using one or a fleet of mobile sensing platforms can be beneficial in filling such data gaps, *e.g.*, in traffic monitoring at temporarily deployed road infrastructure or safety-related sensing at intersections. Insights from this thesis suggest that designing, deploying, and operating such sensing system requires careful consideration of configuration options that align with specific performance metrics. This work hints at many interesting use cases where architectural and data strategies, *i.e.*, determining the minimum amount of data required for specific tasks, still need to be determined. Practice should not be dependent on the IAS industry setting standards maximizing their utility functions, as collecting data at a local scale, even using more traditional MSPs like sensor-equipped cars, buses, trash-trucks, drone, bicycles, and so on, can already offer significant value both at micro and macroscopic level. This is not a blind promotion of the presented sensing paradigm; rather, we do emphasize the importance of readily available data sources in meeting essential needs across various use cases. The algorithmic contributions of this work, together with their application in a traffic control context, imply that already simple algorithmic solutions can be of interest for specific use cases in the real-world.

### Socio-technical understanding of stakeholders

Data and compute are the two most important elements for the deployment and operation of such sensing infrastructure. A brief look at today's digital technology industry shows a concentration of both these resources in the hands of a few corporations, making decisions that reinforce their market dominance, while exerting considerable influence on science and policy and society as a whole. Decision-makers must thus understand the politics and economics of computational infrastructure's impact, drawing insights from other infrastructural shifts, to ensure informed decision-making<sup>3</sup>. Policymakers and practitioners must actively participate in discussions about design, deployment and operation of sensing infrastructure associated with IAS or MSP — for instance, advocating to use the generated data for advanced traffic surveillance — as it is easy to see how missing involvement could lead to a situation of limited influence over data usage. After the deployment of a new digital infrastructure, it's often incredibly challenging, if not entirely undesirable by powerful stakeholders, to reverse or undo it. This rings especially true in the case of a surveillance system fueled by vast quantities of data, as history has demonstrated. Policy-makers and practitioners are sometimes offered to tap into live vehicle operation data provided by tech firms, presented as an attractive alternative to traditional data collection methods — or as only data option. This underscores the shift towards integrating operational insights from emerging IAS or MSP directly into policy frameworks and strategic traffic engineering and planning.

---

<sup>3</sup>Consider again the following scenario; if traffic engineering were solely influenced by CAV companies advocating for pedestrian-free streets to enhance their systems' efficiency and safety, and thus maximize their primary utility function, it could lead to decisions that prioritize corporate interests over public welfare. This illustrates the necessity for proactive engagement by policymakers and transportation practitioners as to ensure these new sensing infrastructures serve the public interest and align with societal values, rather than just serving commercial agendas.

### **Actions for privacy and security**

This sensing paradigm generates unprecedented amounts of data in and of urban environments, underscoring the importance for incorporating comprehensive privacy, security, and ethical frameworks to become integral to the design process, ensuring a responsible approach to urban innovation. Policy makers and practitioners must thus not shy away of modelling and analysing potential risks and harms early in the development process, as to proactively implement measures to mitigate those. Policies must be investigated that may regulate the use of secondary data as ways to preserve public privacy Bloom et al. (2017). Centralizing data, for instance, raises serious privacy and security concerns, risking unauthorized surveillance, profiling, and discrimination among others. Applying privacy mechanisms, decentralization, or federated systems can offer viable solutions and need to be further investigated and developed.

## **6.4 Recommendations for future research**

Although the research conducted in this dissertation has valuable contributions, it serves as an initial step towards cultivating a comprehension of the forthcoming sensing paradigm. The generation of detailed traffic information from IAS or MSP opens up several exciting avenues for traffic researchers to explore the dynamics of cyclists and pedestrians within the urban environments. In the following, we provide further details about prospective avenues for research.

### **Data fusion, multi object tracking and scene understanding**

Future research in the field should aim at harnessing network resources more efficiently by leveraging a deeper understanding of real-world traffic dynamics and the environment. A natural first step is to extend the tracking framework presented in Chapter 3. We propose advancing the model to incorporate more sophisticated kinematic models that could improve short-term prediction accuracy, while keeping inherent computational cost in balance. Furthermore, developing methods to handle sparse data is an interesting area of research. To enhance data association logic and prediction performance, one could explore fine-tuning pre-trained models or transfer learning, as well as incorporating more advanced social and behavioural models into a scene understanding framework. Utilizing rich contextual information from diverse data modalities could improve overall model performance. However, such extensions must be carefully evaluated to ensure they align with existing assumptions underlying the proposed framework. Another interesting direction considers the Markov graph transition probabilities, as these could be informed by learned traffic or environmental patterns. Overall, we propose exploring multi-modality as a research direction, building upon our sensor setup from Chapter 2. By integrating camera data, for instance, a scene graph could be build and represents entities and their relationships, such as buses and pedestrians. One possible avenue may thus be combining symbolic and sub-symbolic approaches with recent generative models to leverage richer data representation, improve predictive models, and enhance holistic understanding.

### **Towards next generation traffic surveillance**

In Chapter 2, we analysed the requirements of the sensing paradigm. In Chapter 4, we explored the utility in a traffic signal control scenario. Future research should explore potential applications in various traffic surveillance contexts, where a systematic approach is needed to map objectives to sensing requirements and network characteristics for specific contexts.



*Traffic Impact Assessment* — The new source of traffic information could for instance be used to assess safety risks for active modes at various locations in the network. Future work could assess the impact of newly introduced infrastructure, *e.g.*, bike lanes, pedestrian zones, or traffic calming measures on cyclist and pedestrian behaviour. For instance, compare mobile sensing platforms' data before and after infrastructure changes, potentially even on a lane-level as suggested in Chapter 3, to assess their impact on traffic flows, safety, and overall mode choice. Surrogate safety measures related to active modes (*e.g.*, time to collision) could be measured for uncontrolled intersections and roundabouts in real-time.

*Traffic Monitoring* — Traffic volume, vehicle speeds and cyclist/pedestrian behaviour may, for instance, help infer high-risk areas, which, in-turn, could lead to recommendations for infrastructure improvements and traffic management strategies. Furthermore, future research could build upon recent works using mobile sensor platforms to monitor infrastructure conditions, for instance, mapping properties of sidewalks, bike lanes, and active mode networks in more general (Vogt et al., 2023). Analysing these properties against spatiotemporal traffic densities, velocities and flows can help develop strategies to enhance active transportation options. The selection of mobile sensor platform type is important, as described in Chapter 2. Attributes such as speed, height, or angle can impact data quality and relevance for specific sensing tasks. Extensions of the traffic control use case presented in Chapter 4, may investigate how to detect congestion, conflicts, or safety hazards involving cyclists and pedestrians. This could be of value to provide timely alerts to traffic control entities, or traffic management centres and road users to mitigate potential risks.

*Traffic Control* — Future work should dig into traffic prioritization strategies to enhance safety. A high level of confidence that the conflict zone is clear before clearing is essential. For instance, the model used to identify occlusion in Chapter 4 should be revisited. Location and timing may impact occlusion severity, suggesting a need for investigation into place-dependent occlusion models. Additional sensors (static or moving) may be employed to verify clearance. However, findings of this thesis suggest investigating the impact of various sensing characteristics described in Chapter 2, such as sensor speed, placement, and coverage, on specific safety metrics. Furthermore, abundance of available information in this paradigm opens up new avenues for exploring self-learning algorithms in traffic control. Building on the idea and concepts presented in the cyclist prioritizing traffic control scenario, adopting a self-learning approach may help optimize intersection control (*e.g.*, balancing the delay and stops of cyclists and cars), better predict cyclist trajectories over intersections, and better perform at determining clearance times for traffic conflicts. Moreover, it may be worth investigating the impact of how clustered groups can improve intersection throughput by reducing green extension times.

### **Insights into behaviours, modes and activities**

This new type of data may also enhance our comprehension of pedestrians and cyclist behaviours, modes and activities. The proposed system design presented in Chapter 2 hints to the simultaneous use of internal sensor platform data with external environment perception; recent

work shows how pedestrian activity and transportation modes can be classified using images collected from moving vehicles (Salazar-Miranda et al., 2023). Next steps could analyse factors like interactions with vehicular traffic and compliance with traffic rules could, for instance, guide educational campaigns, and traffic management strategies to promote safer interactions between different road users. The percentage of cyclists with helmets, potentially in relation to speed information, could be measured and similarly lead to strategies for improving helmet use. Furthermore, unmarked crossings points and the start of a cyclist's turning movement could be identified, offering valuable information for reevaluating current infrastructure. However, it is crucial to carefully consider concerns around privacy as presented in the last chapter of this thesis.

### **On the role of privacy engineering**

As automation intersects with human interaction, prioritizing privacy engineering is crucial. We advocate for an interdisciplinary approach that embeds privacy considerations by design. For example, the centralized tracking scheme proposed in Chapter 3 direct attention to latent privacy concerns. Future work investigating the impact on target privacy and application utility could provide valuable insights when designing more decentralized or federated system. Furthermore, future research should put to the test existing privacy frameworks and location privacy preserving mechanisms and evaluate their appropriateness theoretically and quantitatively, within specific traffic engineering applications. Chapter 5 provided theoretical indications that existing LPPMs, initially developed for user-centric LBS, may fail in providing both utility necessary for the functioning of safety-relevant applications of the system and privacy for the non-users. We suggest further investigation into data minimization strategies, as they hold promise for both tracking individual entities (as discussed in Chapters 3 and 4) and providing alternative solutions that support privacy-enhancing systems through observations confirming the absence of targets. Overall, researchers should elaborate on the role and impact of this new mobile surveillance infrastructure on society and the urban space. Normalizing surveillance, as is currently happening at scale, *e.g.*, with video surveillance from doorbells, erodes a right to privacy and goes a slippery slope with increasing deployments of IAS in urban environments where the capture of human "objects" is justified by safety.

# Bibliography

- Ahmed, T., A. Kapadia, V. Potluri, M. Swaminathan (2018) Up to a Limit? Privacy Concerns of Bystanders and Their Willingness to Share Additional Information with Visually Impaired Users of Assistive Technologies, *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, 2(3), pp. 1–27.
- Akter, T., T. Ahmed, A. Kapadia, S. M. Swaminathan (2020) Privacy considerations of the visually impaired with camera based assistive technologies: Misrepresentation, impropriety, and fairness, in: *Proceedings of the 22nd International ACM SIGACCESS Conference on Computers and Accessibility*, ASSETS '20, ACM, pp. 1–14.
- Akyildiz, I. F., W. Su, Y. Sankarasubramaniam, E. Cayirci (2002) A survey on sensor networks, *IEEE Communications Magazine*, 40(8), pp. 102–114.
- Alessandretti, L., U. Aslak, S. Lehmann (2020) The scales of human mobility, *Nature*, 587(7834), pp. 402–407.
- Allen, H. (1982) You are where you live, *Washington Post*.
- Andrés, M., N. Bordenabe, K. Chatzikokolakis, C. Palamidessi (2013) Geo-indistinguishability: differential privacy for location-based systems, in: *Proceedings of the 2013 ACM SIGSAC conference on Computer & communications security*, CCS '13, ACM, pp. 901–914.
- Anjomshoaa, A., F. Duarte, D. Rennings, T. J. Matarazzo, P. deSouza, C. Ratti (2018) City Scanner: Building and Scheduling a Mobile Sensing Platform for Smart City Services, *IEEE Internet of Things Journal*, 5(6), pp. 4567–4579.
- Apple (2020) Apple and google partner on COVID-19 contact tracing technology, <https://www.apple.com/newsroom/2020/04/apple-and-google-partner-on-covid-19-contact-tracing-technology/>.
- Aquino, S. (2024) The new waymo-hyundai partnership reminds that autonomous vehicles matter for accessibility, agency, <https://www.forbes.com/sites/stevenaquino/2024/10/04/the-new-waymo-hyundai-partnership-reminds-that-autonomous-vehicles-matter-for-accessibility-a-gency/>.
- Arellana, J., M. Saltarín, A. Margarita Larrañaga, C. A. González, Virginia I. Henao (2020) Developing an urban bikeability index for different types of cyclists as a tool to prioritise bicycle infrastructure investments, *Transportation Research Part A: Policy and Practice*, 139, pp. 310–334.

- Asuquo, P., H. Cruickshank, J. Morley, C. P. A. Ogah, A. Lei, W. Hathal, S. Bao, Z. Sun (2018) Security and Privacy in Location-Based Services for Vehicular and Mobile Communications: An Overview, Challenges, and Countermeasures, *IEEE Internet of Things Journal*, 5(6), pp. 4778–4802.
- Bar-Shalom, Y., S. Blackman, R. J. Fitzgerald (2007) Dimensionless score function for multiple hypothesis tracking, *IEEE Transactions on Aerospace and Electronic Systems*, 43(1), pp. 392–400.
- Bar-Shalom, Y., X. R. Li (1993) *Estimation and Tracking: Principles, Techniques, and Software*, Artech House, Inc.
- Barthélemy, J., N. Verstaevael, H. Forehead, P. Perez (2019) Edge-Computing Video Analytics for Real-Time Traffic Monitoring in a Smart City, *Sensors*, 19(9).
- Batkovic, I., M. Zanon, N. Lubbe, P. Falcone (2018) A computationally efficient model for pedestrian motion prediction, in: *2018 European Control Conference (ECC)*, pp. 374–379.
- Bertsekas, D. P. (1988) The auction algorithm: A distributed relaxation method for the assignment problem, *Annals of operations research*, 14(1), pp. 105–123.
- Bezemer, L. (2020) Rotterdam zet camera-auto's in om corona-overtredingen tegen te gaan, <https://www.ad.nl/rotterdam/rotterdam-zet-camera-auto-s-in-om-corona-overtredingen-tegen-te-gaan~a72c18ab/>.
- Blackman, S. (2004) Multiple hypothesis tracking for multiple target tracking, *IEEE Aerospace and Electronic Systems Magazine*, 19(1), pp. 5–18.
- Blackman, S., R. Popoli (1999) *Design and analysis of modern tracking systems*, Artech House radar library, Artech House, Inc.
- Bliss, L. (2018) Behind the backlash over sidewalk labs' smart city, <https://www.bloomberg.com/news/articles/2018-09-07/what-s-behind-the-backlash-over-sidewalk-labs-smart-city>.
- Bloom, C., J. Tan, J. Ramjohn, L. Bauer (2017) Self-driving cars and data collection: privacy perceptions of networked autonomous vehicles, in: *Proceedings of the Thirteenth USENIX Conference on Usable Privacy and Security*, SOUPS '17, pp. 357–375.
- Bogost, I. (2016) Will Robocars Kick Humans Off City Streets?, *The Atlantic*.
- Boukoros, S., M. Humbert, S. Katzenbeisser, C. Troncoso (2019) On (the lack of) location privacy in crowdsourcing applications, in: *Proceedings of the 28th USENIX Conference on Security Symposium*, USENIX Association, pp. 1859–1876.
- Calabrese, F., F. C. Pereira, G. Di Lorenzo, L. Liu, C. Ratti (2010) The Geography of Taste: Analyzing Cell-Phone Mobility and Social Events, in: Floréen, P., A. Krüger, M. Spasojevic, eds., *Pervasive Computing*, Lecture Notes in Computer Science, Springer Berlin Heidelberg, pp. 22–37.
- Calacci, D., A. Berke, K. Larson, Alex, Pentland (2019) The tradeoff between the utility and risk of location data and implications for public good, *arXiv:1905.09350 [cs, math]*.

- Chatzimilioudis, G., A. Konstantinidis, C. Laoudias, D. Zeinalipour-Yazti (2012) Crowdsourcing with smartphones, *IEEE Internet Computing*, 16(5), pp. 36–44.
- Chen, C., J. Ma, Y. Susilo, Y. Liu, M. Wang (2016) The promises of big data and small data for travel behavior (aka human mobility) analysis, *Transportation research. Part C, Emerging technologies*, 68, pp. 285–299.
- Cheng, Z., J. Caverlee, K. Lee, D. Z. Sui (2011) Exploring millions of footprints in location sharing services, in: *Proceedings of the International AAAI Conference on Web and Social Media*, pp. 81–88.
- Chong, C.-Y., S. P. Kumar (2003) Sensor networks: evolution, opportunities, and challenges, *Proceedings of the IEEE*, 91(8), pp. 1247–1256.
- Clark, D. D., J. Wroclawski, K. R. Sollins, R. Braden (2002) Tussle in cyberspace: defining tomorrow's internet, *SIGCOMM Comput. Commun. Rev.*, 32(4), pp. 347–356.
- Claus, S. (2022) Groningen werkt aan een groene golf voor fietsers, <https://www.trouw.nl/duurzaamheid-economie/fietsers-in-groningen-krijgen-eindelijk-hun-beloofde-groene-golf~ba32495d>.
- Collins, J. B., J. K. Uhlmann (1992) Efficient gating in data association with multivariate Gaussian distributed states, *IEEE Transactions on Aerospace and Electronic Systems*, 28(3), pp. 909–916.
- Cox, I. J., S. L. Hingorani (1996) An efficient implementation of reid's multiple hypothesis tracking algorithm and its evaluation for the purpose of visual tracking, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18(2), pp. 138–150.
- Cox, J. (2022) Data broker is selling location data of people who visit abortion clinics, <https://www.vice.com/en/article/m7vzjb/location-data-abortion-clinics-safegraph-planned-parenthood>.
- Crow (2022) Fietsketting, <https://www.fietsberaad.nl/Kennisbank/Proef-met-fietsketting-geeft-groepen-fietsers-extr>.
- Cui, Q., N. Wang, M. Haenggi (2017) Spatial Point Process Modeling of Vehicles in Large and Small Cities, in: *GLOBECOM 2017 - 2017 IEEE Global Communications Conference*, pp. 1–7.
- D. Göhring, M. S., M. Wang, T. Ganjineh (2011) Radar/lidar sensor fusion for car-following on highways, in: *The 5th International Conference on Automation, Robotics and Applications*, pp. 407–412.
- Daamen, W., J. Heuvel, D. Ton, S. P. Hoogendoorn (2015) Using bluetooth and wifi to unravel real-world slow mode activity travel behaviour, *Presented at 14th International Conference on Travel Behaviour Research*.
- Daamen, W., E. Kinkel, D. Duives, S. P. Hoogendoorn (2017) Monitoring visitor flow and behavior during a festival: Mysteryland case study, in: *Proceedings of the 96th Annual Meeting of the Transportation Research Board*.

- Dabiri, A., A. Hegyi (2018) Personalised optimal speed advice to cyclists approaching an intersection with uncertain green time, *European Control Conference*, ECC 2018, pp. 1666–1671.
- Dabiri, A., A. Hegyi, S. Hoogendoorn (2022) Optimized speed trajectories for cyclists, based on personal preferences and traffic light information - a stochastic dynamic programming approach, *IEEE Transactions on Intelligent Transportation Systems*, 23(2), pp. 777–793.
- Danalet, A., B. Farooq, M. Bierlaire (2014) A Bayesian approach to detect pedestrian destination-sequences from WiFi signatures, *Transportation Research Part C: Emerging Technologies*, 44, pp. 146–170.
- de Montjoye, Y.-A., C. A. Hidalgo, M. Verleysen, V. D. Blondel (2013) Unique in the Crowd: The privacy bounds of human mobility, *Scientific Reports*, 3.
- de Sio, F. S., G. Mecacci, S. Calvert, D. Heikoop, M. Hagenzieker, B. van Arem (2023) Realising Meaningful Human Control Over Automated Driving Systems: A Multidisciplinary Approach, *Minds and Machines*, 33(4), pp. 587–611.
- Denning, T., Z. Dehlawi, T. Kohno (2014) In situ with bystanders of augmented reality glasses: perspectives on recording and privacy-mediating technologies, in: *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, ACM, pp. 2377–2386.
- Ding, A. Y., M. Janssen (2018) Opportunities for applications using 5g networks: Requirements, challenges, and outlook, in: *Proceedings of the 7th International Conference on Telecommunications and Remote Sensing*, ICTRS '18, ACM, pp. 27–34.
- Dollar, P., C. Wojek, B. Schiele, P. Perona (2012) Pedestrian Detection: An Evaluation of the State of the Art, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(4), pp. 743–761.
- Duives, D. C., W. Daamen, S. P. Hoogendoorn (2013) State-of-the-art crowd motion simulation models, *Transportation Research Part C: Emerging Technologies*, 37, pp. 193–209.
- Duives, D. C., T. van Oijen, S. P. Hoogendoorn (2020) Enhancing Crowd Monitoring System Functionality through Data Fusion: Estimating Flow Rate from Wi-Fi Traces and Automated Counting System Data, *Sensors*, 20(21).
- Duives, D. C., G. Wang, J. Kim (2019) Forecasting pedestrian movements using recurrent neural networks: An application of crowd monitoring data, *Sensors*, 19(2).
- Eckhoff, D., C. Sommer (2014) Driving for Big Data? Privacy Concerns in Vehicular Networking, *IEEE Security Privacy*, 12(1), pp. 77–79.
- Enzweiler, M., D. M. Gavrila (2009) Monocular Pedestrian Detection: Survey and Experiments, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(12), pp. 2179–2195.
- Fajans, J., M. Curry (2001) Why bicyclists hate stop signs, *UC Berkeley ACCESS Magazine*, 1(18), pp. 27–31.

- Fang, H., N. Tian, Y. Wang, M. Zhou, M. A. Haile (2018) Nonlinear Bayesian estimation: from Kalman filtering to a broader horizon, *IEEE/CAA Journal of Automatica Sinica*, 5(2), pp. 401–417.
- Feng, Y., D. Duives, W. Daamen, S. P. Hoogendoorn (2021) Data collection methods for studying pedestrian behaviour: A systematic review, *Building and Environment*, 187.
- Feng, Y., Z. Xu, H. Farah, B. van Arem (2024) Does another pedestrian matter? a virtual reality study on the interaction between multiple pedestrians and autonomous vehicles in shared space, *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–14.
- Fietsberaad (2012) Groningen-gives-cyclist-the-green-light-more-often, <https://www.fietsberaad.nl/Kennisbank/Groningen-gives-cyclist-the-green-light-more-often?feed=news>.
- Fløan, I., H. Looijen, P. Smit, J. Hiddink (2016) Intelligente verkeers regel installatie, (ivri) – fase 1, deliverable f: itlc architecture, Tech. rep., Imtech, Siemens, Swarco, Vialis.
- Forlano, L. (2019) Cars and Contemporary Communications| Stabilizing/Destabilizing the Driverless City: Speculative Futures and Autonomous Vehicles, *International Journal of Communication*, 13.
- Fugiglando, U., E. Massaro, P. Santi, S. Milardo, K. Abida, R. Stahlmann, F. Netter, C. Ratti (2019) Driving Behavior Analysis through CAN Bus Data in an Uncontrolled Environment, *IEEE Transactions on Intelligent Transportation Systems*, 20(2), pp. 737–748.
- Furth, P. G., T. H. Muller (1999) Trafcod : A method for stream-based control of actuated traffic signals, in: *Proceedings of the 78th Annual Meeting of the Transportation Research Board*, p. 10–14.
- Furth, P. G., T. H. J. Muller, M. Salomons, T. A. Bertulis, P. J. Koonce (2012) Barrier-free ring structures and pedestrian overlaps in signalized intersection control, *Transportation Research Record: Journal of the Transportation Research Board*, 2311, pp. 132–141.
- Gambs, S., M.-O. Killijian, M. Núñez del Prado Cortez (2014) De-anonymization attack on geolocated data, *Journal of Computer and System Sciences*, 80(8), pp. 1597–1614.
- García-Fernández, Á. F., J. L. Williams, K. Granström, L. Svensson (2018) Poisson Multi-Bernoulli Mixture Filter: Direct Derivation and Implementation, *IEEE Transactions on Aerospace and Electronic Systems*, 54(4), pp. 1883–1901.
- Gavriilidou, A., W. Daamen, Y. Yuan, S. P. Hoogendoorn (2019a) Modelling cyclist queue formation using a two-layer framework for operational cycling behaviour, *Transportation Research Part C: Emerging Technologies*, 105, pp. 468–484.
- Gavriilidou, A., M. J. Wierbos, W. Daamen, Y. Yuan, V. L. Knoop, S. P. Hoogendoorn (2019b) Large-scale bicycle flow experiment: Setup and implementation, *Transportation Research Record*, 2673(5), pp. 709–719.
- Godefrooij, H. (2016) Capaciteit fietspaden bij vri's. data-analyse, richtlijnen en handvatten, Tech. rep., CROW, DTV.

- Gong, V. X., W. Daamen, A. Bozzon, S. P. Hoogendoorn (2020) Crowd characterization for crowd management using social media data in city events, *Travel Behaviour and Society*, 20, pp. 192–212.
- Gonzalez, M. C., C. A. Hidalgo, A.-L. Barabasi (2008) Understanding individual human mobility patterns, *Nature*, 453(7196), pp. 779–782.
- Gordon, A. (2022) San francisco police are using driverless cars as mobile surveillance cameras, <https://www.vice.com/en/article/v7dw8x/san-francisco-police-are-using-driverless-cars-as-mobile-surveillance-cameras>.
- Grauerholz, M. (2017) The city of boston uses data to improve quality of life for residents, *U.S. Green Building Council*.
- Greenberg, A. (2019) This tesla mod turns a model s into a mobile 'surveillance station, <https://www.wired.com/story/tesla-surveillance-detection-scout/>.
- Gruteser, M., B. Hoh (2005) On the Anonymity of Periodic Location Samples, in: *Security in Pervasive Computing*, Lecture Notes in Computer Science, Springer, pp. 179–192.
- Guariglia, M. (2023) The impending privacy threat of self-driving cars, <https://www.eff.org/deeplinks/2023/08/impending-privacy-threat-self-driving-cars>.
- Han, B., V. Gopalakrishnan, L. Ji, S. Lee (2015) Network function virtualization: Challenges and opportunities for innovations, *IEEE Communications Magazine*, 53(2), pp. 90–97.
- Harris, M. (2022a) The Radical Scope of Tesla's Data Hoard, <https://spectrum.ieee.org/tesla-autopilot-data-scope>.
- Harris, M. (2022b) Tesla's autopilot depends on a deluge of data, <https://spectrum.ieee.org/tesla-autopilot-data-deluge>.
- Hawkins, A. (2019) Waymo is bringing its self-driving cars to map los angeles, <https://www.theverge.com/2019/10/7/20903034/waymo-self-driving-test-la-mds-data>.
- Hawkins, A. (2023) Waymo has 7.1 million driverless miles — how does its driving compare to humans?, <https://www.theverge.com/2023/12/20/24006712/waymo-driverless-million-mile-safety-compare-human>.
- Heathman, A. (2017) Protect your privacy with anti-surveillance clothing, <https://www.wired.co.uk/article/anti-surveillance-clothing-adam-harvey>.
- Hoogendoorn, S. P., P. H. L. Bovy (2004) Pedestrian route-choice and activity scheduling theory and models, *Transportation Research Part B: Methodological*, 38(2), pp. 169–190.
- Hoyle, R., R. Templeman, S. Armes, D. Anthony, D. Crandall, A. Kapadia (2014) Privacy behaviors of lifeloggers using wearable cameras, in: *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, ACM, pp. 571–582.
- Hull, B., V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan, S. Madden (2006) CarTel: A distributed mobile sensor computing system, in: *Proc. ACM SenSys*, pp. 125–138.



- Jackson, M., P. Zave (1995) Deriving specifications from requirements: an example, in: *Proceedings of the 17th international conference on Software engineering*, ICSE '95, ACM, pp. 15–24.
- Jacobs, K. (2022) Toronto wants to kill the smart city forever, <https://www.technologyreview.com/2022/06/29/1054005/toronto-kill-the-smart-city/>.
- Jayawardana, V., C. Wu (2022) Learning eco-driving strategies at signalized intersections, in: *2022 European Control Conference (ECC)*, pp. 383–390.
- Jazwinski, A. H. (1970) *Stochastic Processes and Filtering Theory*, vol. 64 of *Mathematics in Science and Engineering*, Elsevier.
- Jonker, R., A. Volgenant (1987) A shortest augmenting path algorithm for dense and sparse linear assignment problems, *Computing*, 38(4), pp. 325–340.
- Julier, S. J., J. K. Uhlmann (1997) New extension of the Kalman filter to nonlinear systems, in: *Signal Processing, Sensor Fusion, and Target Recognition VI*, vol. 3068, International Society for Optics and Photonics, SPIE, pp. 182–193.
- Jurdak, R., K. Zhao, J. Liu, M. AbouJaoude, M. Cameron, D. Newth (2015) Understanding Human Mobility from Twitter, *PLOS ONE*, 10(7).
- Kadota, I., A. Sinha, E. Uysal-Biyikoglu, R. Singh, E. Modiano (2018) Scheduling Policies for Minimizing Age of Information in Broadcast Wireless Networks, *arXiv:1801.01803 [cs, math]*.
- Kalman, R. E. (1960) A new approach to linear filtering and prediction problems, *Journal of basic Engineering*, 82(1), pp. 35–45.
- Keller, C. G., D. M. Gavrila (2014) Will the pedestrian cross? a study on pedestrian path prediction, *IEEE Transactions on Intelligent Transportation Systems*, 15(2), pp. 494–506.
- Kim, C., F. Li, A. Ciptadi, J. M. Rehg (2015) Multiple hypothesis tracking revisited, in: *2015 IEEE International Conference on Computer Vision (ICCV)*, pp. 4696–4704.
- Kirubarajan, T., Y. Bar-Shalom, K. Pattipati, I. Kadar (2000) Ground target tracking with variable structure IMM estimator, *IEEE Transactions on Aerospace and Electronic Systems*, 36(1), pp. 26–46.
- Klein, D. (1989) You are where you live - neighborhoods by the numbers - how zip codes can pigeonhole people like widgets in a digit factory, *Los Angeles Times*.
- Klein, L. A., M. K. Mills, D. Gibson (2006) Traffic detector handbook : third edition. volume ii, Tech. Rep. FHWA-HRT-06-139, FHWA.
- Knoop, V. L., P. J. Buist, C. C. J. M. Tiberius, B. v. Arem (2012) Automated lane identification using precise point positioning an affordable and accurate GPS technique, in: *2012 15th International IEEE Conference on Intelligent Transportation Systems*, pp. 939–944.

- Kooij, J. F. P., F. Flohr, E. A. I. Pool, D. M. Gavrilu (2019) Context-Based Path Prediction for Targets with Switching Dynamics, *International Journal of Computer Vision*, 127(3), pp. 239–262.
- Kreutz, D., F. M. V. Ramos, P. E. Veríssimo, C. E. Rothenberg, S. Azodolmolky, S. Uhlig (2015) Software-defined networking: A comprehensive survey, *Proceedings of the IEEE*, 103(1), pp. 14–76.
- Krishanth, K., X. Chen, R. Tharmarasa, T. Kirubarajan, M. McDonald (2017) The Social Force PHD Filter for Tracking Pedestrians, *IEEE Transactions on Aerospace and Electronic Systems*, 53(4), pp. 2045–2059.
- Krishanth, K., R. Tharmarasa, T. Kirubarajan, P. Valin, E. Meger (2014) Prediction and retrodiction algorithms for path-constrained targets, *IEEE Transactions on Aerospace and Electronic Systems*, 50(4), pp. 2746–2761.
- Krumm, J. (2007) Inference Attacks on Location Tracks, in: *Pervasive Computing*, Lecture Notes in Computer Science, Springer, pp. 127–143.
- Krumm, J. (2009) A survey of computational location privacy, *Personal and Ubiquitous Computing*, 13(6), pp. 391–399.
- Kuhn, H. W. (1955) The hungarian method for the assignment problem, *Naval Research Logistics Quarterly*, 2(1-2), pp. 83–97.
- Kurien, T. (1990) Issues in the design of practical multitarget tracking algorithms, in: *Multitarget-Multisensor Tracking: Advanced Applications*, Artech House, pp. 43–83.
- Laaroussi, Z., R. Morabito, T. Taleb (2018) Service provisioning in vehicular networks through edge and cloud: An empirical analysis, in: *2018 IEEE Conference on Standards for Communications and Networking (CSCN)*, pp. 1–6.
- Lai, M. (2021) *Intersection control for cyclists with iSignum*, Bachelor thesis, Technische Universiteit Delft, <https://repository.tudelft.nl/record/uuid:b29cd450-cefd-4f38-a5eb-04ccf4da78f4>.
- Lang, N., A. Herrmann, M. Hagenmaier, M. Richter (2020) Can self-driving cars stop the urban mobility meltdown?, <https://www.bcg.com/publications/2020/how-autonomous-vehicles-can-benefit-urban-mobility>.
- Lee, O., A. Rasch, A. L. Schwab, M. Dozza (2020) Modelling cyclists' comfort zones from obstacle avoidance manoeuvres, *Accident Analysis & Prevention*, 144.
- Lee, U., B. Zhou, M. Gerla, E. Magistretti, P. Bellavista, A. Corradi (2006) Mobeyes: smart mobs for urban monitoring with a vehicular sensor network, *IEEE Wireless Communications*, 13(5), pp. 52–57.
- Leonardi, P. M. (2010) From road to lab to math: The co-evolution of technological, regulatory, and organizational innovations for automotive crash testing, *Social Studies of Science*, 40(2), pp. 243–274.

- Liu, B., P. Brass, O. Dousse, P. Nain, D. Towsley (2005) Mobility Improves Coverage of Sensor Networks, in: *Proceedings of the 6th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, MobiHoc '05, ACM, pp. 300–308.
- Liu, L., M. Gruteser (2021) Edgesharing: Edge assisted real-time localization and object sharing in urban streets, in: *IEEE INFOCOM 2021 - IEEE Conference on Computer Communications*, pp. 1–10.
- López-Araquistain, J., A. J. Jarama, J. A. Besada, G. d. Miguel, J. R. Casar (2019) A new approach to map-assisted Bayesian tracking filtering, *Information Fusion*, 45, pp. 79–95.
- Luber, M., G. D. Tipaldi, K. O. Arras (2011) Place-dependent people tracking, *The International Journal of Robotics Research*, 30(3), pp. 280–293.
- Lusk, A. (2016) Promoting bicycling through creative design: innovations for bicycles and cycling facilities, <https://cdn1.sph.harvard.edu/wp-content/uploads/sites/1666/2016/07/Promoting-Bicycling-through-Creative-Design.pdf>.
- Ly, M. V., S. Martin, M. M. Trivedi (2013) Driver classification and driving style recognition using inertial sensors, in: *2013 IEEE Intelligent Vehicles Symposium (IV)*, pp. 1040–1045.
- Malaker, T., Q. Meng (2024) Urban disparity analytics using gis: A systematic review, *Sustainability*, 16(14).
- Massaro, E., C. Ahn, C. Ratti, P. Santi, R. Stahlmann, A. Lamprecht, M. Roehder, M. Huber (2017) The Car as an Ambient Sensing Platform [Point of View], *Proceedings of the IEEE*, 105(1), pp. 3–7.
- Matarazzo, T. J., P. Santi, S. N. Pakzad, K. Carter, C. Ratti, B. Moaveni, C. Osgood, N. Jacob (2018) Crowdsensing Framework for Monitoring Bridge Vibrations Using Moving Smartphones, *Proceedings of the IEEE*, 106(4), pp. 577–593.
- Meel, E. (2013) *Red light running by cyclists*, Master thesis, Technische Universiteit Delft.
- Milakis, D., M. Kroesen, B. van Wee (2018) Implications of automated vehicles for accessibility and location choices: Evidence from an expert-based experiment, *Journal of Transport Geography*, 68, pp. 142–148.
- Milakis, D., B. van Arem, B. van Wee (2017) Policy and society related implications of automated driving: A review of literature and directions for future research, *Journal of Intelligent Transportation Systems*, 21(4), pp. 324–348.
- Minet, L., R. Liu, M.-F. Valois, J. Xu, S. Weichenthal, M. Hatzopoulou (2018) Development and Comparison of Air Pollution Exposure Surfaces Derived from On-Road Mobile Monitoring and Short-Term Stationary Sidewalk Measurements, *Environmental Science & Technology*, 52(6), pp. 3512–3519.
- Mladenovic, M. N., T. McPherson (2016) Engineering social justice into traffic control for self-driving vehicles?, *Science and engineering ethics*, 22, pp. 1131–1149.

- Mohan, P., V. N. Padmanabhan, R. Ramjee (2008) Nericell: rich monitoring of road and traffic conditions using mobile smartphones, in: *Proceedings of the 6th ACM conference on Embedded network sensor systems*, SenSys '08, ACM, pp. 323–336.
- Morabito, R., V. Cozzolino, A. Y. Ding, N. Beijar, J. Ott (2018) Consolidate iot edge computing with lightweight virtualization, *IEEE Network*, 32(1), pp. 102–111.
- Morabito, R., R. Petrolo, V. Loscrì, N. Mitton, G. Ruggeri, A. Molinaro (2017) Lightweight virtualization as enabling technology for future smart cars, in: *2017 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*, pp. 1238–1245.
- Murty, K. G. (1968) An algorithm for ranking all the assignments in order of increasing cost, *Operations Research*, 16(3), pp. 682–687.
- Nguyen, T.-H., T. G. Vu, H.-L. Tran, K.-S. Wong (2022) Emerging privacy and trust issues for autonomous vehicle systems, in: *2022 International Conference on Information Networking*, pp. 52–57.
- Ohn-Bar, E., M. M. Trivedi (2016) Looking at Humans in the Age of Self-Driving and Highly Automated Vehicles, *IEEE Transactions on Intelligent Vehicles*, 1, pp. 90–104.
- O’Keeffe, K. P., A. Anjomshoaa, S. H. Strogatz, P. Santi, C. Ratti (2019) Quantifying the sensing power of vehicle fleets, *Proceedings of the National Academy of Sciences*, 116(26), pp. 12752–12757.
- Olofsson, J., E. Brekke, T. I. Fossen, T. A. Johansen (2017) Spatially indexed clustering for scalable tracking of remotely sensed drift ice, in: *2017 IEEE Aerospace Conference*, pp. 1–13.
- Olteanu, A.-M., K. Huguenin, I. Dacosta, J.-P. Hubaux (2018) Consensual and privacy-preserving sharing of multi-subject and interdependent data, 25th Network and Distributed System Security Symposium (NDSS).
- Pannetier, B., K. Benameur, V. Nimier, M. Rombaut (2004) Ground moving target tracking with road constraint, in: *Signal Processing, Sensor Fusion, and Target Recognition XIII*, vol. 5429, pp. 138–149.
- Parikka, J. (2023) *Operational Images: From the Visual to the Invisual*, University of Minnesota Press.
- Pashkevich, A. (2013) *Automated Analysis Methods for the Assessment of Bicycle Infrastructure and Demand*, Ph.D. thesis, Università di Bologna.
- Phillips, D., M. Curry (2002) Privacy and the phenetic urge: Geodemographics and the changing spatiality of local practice, in: *Surveillance as Social Sorting: Privacy, Risk, and Digital Discrimination*, Routledge, p. 137–152.
- Plantin, J.-C. (2018) Digital Traces in Context| Google Maps as Cartographic Infrastructure: From Participatory Mapmaking to Database Maintenance, *International Journal of Communication*, 12(0).

- Poelman, M., A. Hegyi, A. Verbraeck, J. van Lint (2020) Sensitivity analysis to define guidelines for predictive control design, *Transportation Research Record*, 2674(6), pp. 385–398.
- Primault, V., A. Boutet, S. B. Mokhtar, L. Brunie (2019) The Long Road to Computational Location Privacy: A Survey, *IEEE Communications Surveys Tutorials*, 21(3), pp. 2772–2793.
- PTV (2016) Vissim 11 traffic simulation software, <https://www.ptvgroup.com/en/products/ptv-vissim>, PTV Group, Karlsruhe, Germany.
- Pyrgelis, A., C. Troncoso, E. De Cristofaro (2017) What does the crowd say about you? evaluating aggregation-based location privacy, *Proceedings on Privacy Enhancing Technologies*, pp. 156–176.
- Qiu, H., F. Ahmad, F. Bai, M. Gruteser, R. Govindan (2018) AVR: Augmented Vehicular Reality, in: *Proceedings of the 16th Annual International Conference on Mobile Systems, Applications, and Services*, MobiSys '18, ACM, New York, NY, USA, pp. 81–95.
- Rahmathullah, A. S., Á. F. García-Fernández, L. Svensson (2017) Generalized optimal sub-pattern assignment metric, in: *2017 20th International Conference on Information Fusion (Fusion)*, pp. 1–8.
- Reggiani, G., A. M. Salomons, M. Sterk, Y. Yuan, S. O'Hern, W. Daamen, S. Hoogendoorn (2022) Bicycle network needs, solutions, and data collection systems: A theoretical framework and case studies, *Case Studies on Transport Policy*, 10(2), pp. 927–939.
- Reid, D. (1979) An algorithm for tracking multiple targets, *IEEE Transactions on Automatic Control*, 24(6), pp. 843–854.
- Reuter, S., B.-T. Vo, B.-N. Vo, K. Dietmayer (2014) The Labeled Multi-Bernoulli Filter, *IEEE Transactions on Signal Processing*, 62(12), pp. 3246–3260.
- Ristic, B., S. Arulampalam, N. Gordon (2004) *Beyond the Kalman Filter: Particle Filters for Tracking Applications*, Artech House.
- Rubio, A. I. (2015) Lange files door regensensor voor fietsers, <https://www.ad.nl/rotterdam/lange-files-door-regensensor-voor-fietsers~ada42c0f/>.
- Rudenko, A., L. Palmieri, M. Herman, K. M. Kitani, D. M. Gavrila, K. O. Arras (2020) Human motion trajectory prediction: a survey, *The International Journal of Robotics Research*, 39(8), pp. 895–935.
- Safecast (2022) <https://safecast.org/>.
- Salazar-Miranda, A., F. Zhang, M. Sun, P. Leoni, F. Duarte, C. Ratti (2023) Smart curbs: Measuring street activities in real-time using computer vision, *Landscape and Urban Planning*, 234.
- Sambasivan, N., S. Kapania, H. Highfill, D. Akrong, P. Paritosh, L. M. Aroyo (2021) “everyone wants to do the model work, not the data work”: Data cascades in high-stakes ai, in: *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, CHI '21, ACM, pp. 1–15.

- Sanchez, F., R. Blanco, J. L. Diez (2016) Better Together: Cooperative Technologies Will Be Vital to the Development of Highly Autonomous Vehicles Operating in Complex Urban Environments, *Vision Zero International*.
- Schläpfer, M., L. Dong, K. O’Keeffe, P. Santi, M. Szell, H. Salat, S. Anklesaria, M. Vazifeh, C. Ratti, G. B. West (2021) The universal visitation law of human mobility, *Nature*, 593(7860), pp. 522–527.
- Shokri, R., J. Freudiger, J.-P. Hubaux (2010) A Unified Framework for Location Privacy, *Proceedings 3rd Symposium on Hot Topics in Privacy Enhancing Technologies (HotPETs 2010)*.
- Shoval, N., B. McKercher, E. Ng, A. Birenboim (2011) Hotel location and tourist activity in cities, *Annals of Tourism Research*, 38(4), pp. 1594–1612.
- Song, C., Z. Qu, N. Blumm, A.-L. Barabási (2010) Limits of Predictability in Human Mobility, *Science*, 327(5968), pp. 1018–1021.
- Song, D., R. Tharmarasa, T. Kirubarajan, X. N. Fernando (2018) Multi-Vehicle Tracking With Road Maps and Car-Following Models, *IEEE Transactions on Intelligent Transportation Systems*, 19(5), pp. 1375–1386.
- Soteropoulos, A., M. Berger, F. Ciari (2019) Impacts of automated vehicles on travel behaviour and land use: an international review of modelling studies, *Transport Reviews*, 39(1), pp. 29–49.
- Srinivasan, W. W. V., K.-C. Chua (2007) Trade-offs Between Mobility and Density for Coverage in Wireless Sensor Networks, in: *Proceedings of the 13th Annual ACM International Conference on Mobile Computing and Networking, MobiCom ’07*, ACM, pp. 39–50.
- Stahie, S. (2019) Amazon’s ring surveillance dragnet challenged by digital rights activists, <https://www.bitdefender.com/blog/hotforsecurity/amazons-ring-surveillance-drag-net-challenge-d-digital-rights-activists>.
- Statista (2022) Connected cars as a share of the total car parc in the european union (eu-27) in 2021, with a forecast for 2025, 2030, and 2035, <https://www.statista.com/statistics/1276381/connected-car-fleet-as-a-share-of-total-car-parc-in-european-union/>.
- Strauss, J., L. F. Miranda-Moreno (2017) Speed, travel time and delay for intersections and road segments in the montreal network using cyclist smartphone gps data, *Transportation Research Part D: Transport and Environment*, 57, pp. 155–171.
- Sun, L., Y. Yin (2017) Discovering themes and trends in transportation research using topic modeling, *Transportation Research Part C: Emerging Technologies*, 77, pp. 49–66.
- Takagi, S., Y. Cao, Y. Asano, M. Yoshikawa (2020) Geo-Graph-Indistinguishability: Location Privacy on Road Networks Based on Differential Privacy, *arXiv:2010.13449 [cs]*.
- Tan, D., M. Younis, W. Lalouani, S. Fan, S. Song (2024) A novel pedestrian road crossing simulator for dynamic traffic light scheduling systems, *Journal of Intelligent Transportation Systems*, 28(5), pp. 636–650.

- Tettamanti, T., I. Varga, Z. Szalay (2016) Impacts of autonomous cars from a traffic engineering perspective, *Periodica Polytechnica Transportation Engineering*, 44(4), pp. 244–250.
- Toh, C. K., J. A. Sanguesa, J. C. Cano, F. J. Martinez (2020) Advances in smart roads for future smart cities, *Proceedings of the Royal Society A*, 476(2233).
- Twaddle, H., T. Schendzielorz, O. Fakler (2014) Bicycles in Urban Areas: Review of Existing Methods for Modeling Behavior, *Transportation Research Record*, 2434(1), pp. 140–146.
- Ulmke, M., W. Koch (2006) Road-map assisted ground moving target tracking, *IEEE Transactions on Aerospace and Electronic Systems*, 42(4), pp. 1264–1274.
- van der Vliet, H., P. Luns, J. Zee, J. Douma (2016) ivri fase 1 - deliverable e: Use case beschrijvingen, <https://www.ivera.nl/wp-content/uploads/2016/01/Deliverable-E-Beschrijving-use-cases.pdf>.
- van Dijck, G., E. Greweldinger, J. van Run, M. Bergmans, L. Krol, L. Prinsen, L. Misdom (2018) Onderzoek detectieconfiguratie en signaalgroepafhandeling, <https://www-crow-nl.tu-delft.idm.oclc.org/downloads/pdf/verkeer-en-vervoer/verkeersmanagement/verkeersregelsinstallaties/uta276kll1828-01-iver-detectieonderzoek>.
- van Oijen, T. P., W. Daamen, S. P. Hoogendoorn (2020) Estimation of a recursive link-based logit model and link flows in a sensor equipped network, *Transportation Research Part B: Methodological*, 140, pp. 262–281.
- Verbeeke, R. (2020) *Green wave-apps for cyclists*, Bachelor thesis, Technische Universiteit Delft, <https://repository.tudelft.nl/record/uuid:84a04f0c-818a-45ec-a3b4-b1cc0d2b3dd5>.
- Versichele, M., L. de Groote, M. Claeys Bouuaert, T. Neutens, I. Moerman, N. Van de Weghe (2014) Pattern mining in tourist attraction visits through association rule learning on Bluetooth tracking data: A case study of Ghent, Belgium, *Tourism Management*, 44, pp. 67–81.
- Vial, A., W. Daamen, A. Y. Ding, B. van Arem, S. P. Hoogendoorn (2020) AMSense: How Mobile Sensing Platforms Capture Pedestrian/Cyclist Spatiotemporal Properties in Cities, *IEEE Intelligent Transportation Systems Magazine*, 14(1), pp. 29–43.
- Vial, A., G. Hendebay, W. Daamen, B. van Arem, S. P. Hoogendoorn (2023a) Framework for Network-Constrained Tracking of Cyclists and Pedestrians, *IEEE Transactions on Intelligent Transportation Systems*, 24(3), pp. 3282–3296.
- Vial, A., M. Salomons, W. Daamen, B. van Arem, S. Hoogendoorn-Lanser, S. P. Hoogendoorn (2023b) Prioritizing Cyclists at Signalized Intersections Using Observations from Connected Autonomous Vehicles, *Transportation Research Record*, 2677(12), pp. 29–43.
- Vo, B., W. Ma (2006) The Gaussian Mixture Probability Hypothesis Density Filter, *IEEE Transactions on Signal Processing*, 54(11), pp. 4091–4104.
- Vogt, J., M. Ilic, K. Bogenberger (2023) A mobile mapping solution for vru infrastructure monitoring via low-cost lidar-sensors, *Journal of Location Based Services*, 17(4), pp. 389–411.

- Wagner, I., D. Eckhoff (2018) Technical Privacy Metrics: a Systematic Survey, *ACM Computing Surveys*, 51, pp. 1–38.
- Wang, C., F. Haider, X. Gao, X. You, Y. Yang, D. Yuan, H. M. Aggoune, H. Haas, S. Fletcher, E. Hepsaydir (2014) Cellular architecture and key technologies for 5g wireless communication networks, *IEEE Communications Magazine*, 52(2), pp. 122–130.
- Wang, H., B. Wang, B. Liu, X. Meng, G. Yang (2017) Pedestrian recognition and tracking using 3d LiDAR for autonomous vehicle, *Robotics and Autonomous Systems*, 88, pp. 71–78.
- Webster, F. (1958) Traffic signal settings, *Road Research Technical Paper 39*, London.
- Wierbos, M., V. Knoop, R. Bertini, S. Hoogendoorn (2021) Influencing the queue configuration to increase bicycle jam density and discharge rate: An experimental study on a single path, *Transportation Research Part C: Emerging Technologies*, 122.
- Wirz, M., T. Franke, D. Roggen, E. Mitleton-Kelly, P. Lukowicz, G. Tröster (2012) Inferring Crowd Conditions from Pedestrians' Location Traces for Real-Time Crowd Monitoring during City-Scale Mass Gatherings, in: *2012 IEEE 21st International Workshop on Enabling Technologies: Infrastructure for Collaborative Enterprises*, pp. 367–372.
- Xu, F., Z. Tu, Y. Li, P. Zhang, X. Fu, D. Jin (2017) Trajectory Recovery From Ash: User Privacy Is NOT Preserved in Aggregated Mobility Data, in: *Proceedings of the 26th International Conference on World Wide Web*, pp. 1241–1250.
- Yang, Y., A. Heppenstall, A. Turner, A. Comber (2019) Who, where, why and when? using smart card and social media data to understand urban mobility, *ISPRS International Journal of Geo-Information*, 8(6).
- Zang, H., J. Bolot (2011) Anonymization of location data does not work: a large-scale measurement study, in: *Proceedings of the 17th annual international conference on Mobile computing and networking*, MobiCom '11, ACM, pp. 145–156.
- Zhang, L., H. Han, M. Zhou, Y. Al-Turki, A. Abusorrah (2021a) An Improved Discriminative Model Prediction Approach to Real-Time Tracking of Objects With Camera as Sensors, *IEEE Sensors Journal*, 21(15), pp. 17308–17317.
- Zhang, S., Y. Feng, L. Bauer, L. F. Cranor, A. Das, N. Sadeh (2021b) “Did you know this camera tracks your mood?”: Understanding Privacy Expectations and Preferences in the Age of Video Analytics, *Proceedings on Privacy Enhancing Technologies*, 2021(2).
- Zhao, Y., I. Wagner (2019) On the Strength of Privacy Metrics for Vehicular Communication, *IEEE Transactions on Mobile Computing*, 18(2), pp. 390–403.
- Zheng, J., M. Gao (2018) Tracking Ground Targets with a Road Constraint Using a GMPHD Filter, *Sensors (Basel, Switzerland)*, 18(8), p. 2723.



# Summary

To promote safer, more equitable, and environmentally friendly transportation options, governments are aiming to boost the proportion of trips made walking and cycling. Despite the acknowledged importance of active mode research, however, a relative lack of empirical approaches, advanced tools and high-quality data persists in understanding the dynamics of when, where, and how people walk and cycle. This resulted in relatively limited scientific advancements and mostly timid implementation of these findings in cities.

This thesis leverages a novel sensing paradigm, driven by the deployment of intelligent autonomous systems (IAS) in cities (*e.g.*, self-driving cars, drones, and other instances), which results in an unprecedented surge in the generation of human features — such as their movement patterns, shapes, and contextual information — at great temporal and spatial scale. Importantly, the thesis also applies to other moving sensor platforms (MSP) equipped with advanced sensing capabilities that do not rely on this information to self-operate, such as traditional buses, bicycles, or even humans using bodycams. When autonomous, however, these systems are specifically designed to capture spatiotemporal data of all objects and traffic in their surrounding environment, including pedestrians and cyclists, as they require this information to self-operate. The generated data may represent a valuable new source of information for traffic researchers and practitioners, however, have implications on privacy or broader potential of misuse.

With this in mind, the objective of this thesis is *to investigate the potential utility and societal impact of exploiting pedestrian and cyclist data from moving sensors (IAS or MSP) in the context of traffic research and practice*. The use of the term *exploitation* is twofold: 1) the development of methodological knowledge and tools to extract traffic information from the data, and 2) under what conditions (*e.g.*, limitations, risks, and harms) it contributes to technological, societal and economic development.

To achieve this goal and gain a holistic understanding of this emerging paradigm, this thesis builds upon an interdisciplinary research approach, bridging transport research with other relevant scientific communities. To establish a solid foundation for our assumptions and design choices moving forward, we derived a theoretical framework by synthesizing diverse data sources and expert knowledge. This conceptualization, tailored specifically for active mode research and practice, focuses on key components of the system as sensing, networking, processing, and communication units, enabled us to discern requirements, along with functional challenges, and to outline research to be performed with the generated data (Chapter 2). We then refine our analysis by examining both the temporal and spatial dimensions of the data, as it could be generated at different levels of granularity and accuracy. We derived a methodology capable of consistently yielding updated active mode location information throughout the entire network, that allows us to reconstruct individual trajectories across a road network, using mo-

bile spatially distributed sensor platforms (Chapter 3). Using this approach, and leveraging the insights developed so far, we investigate and demonstrate how this new type of data can provide valuable insights and help optimize traffic control and management strategies. We defined a traffic control scenario, and used modelling and simulation to generate scenario-specific data at scale (Chapter 4). In fact, due to the lack of adequate real-world data, simulation was used when necessary to address the research questions 2, 3. When gathering data on individuals, however, there is an inherent trade-off between utility and privacy. The sensing paradigm's broader socio-technical impact is examined by comparing it with current systems, emphasizing location privacy implications for non-users (pedestrians and cyclists) (Chapter 5). The following paragraphs detail the findings of this research with respect to these four topics.

### **Towards new system designs for active mode mobile sensing**

We present a theoretical design for a novel mobile sensing system (AMSense) that uses vehicles as mobile sensing nodes in a network to capture spatiotemporal properties of pedestrians and cyclists in urban environments. In this dynamic, multi-sensor approach, real-time data, algorithms, and models are fused to estimate presence, positions and movements of active modes with information generated by a fleet of mobile sensing platforms. It represents a scalable system that provides answers to spatiotemporal resolution, intrusiveness, and dynamic network conditions. The moving platform may thus increase observational robustness through data from multiple-sources. In this chapter, we motivate the need and show the potential of such a sensing paradigm, which supports a host of new research and application development. We investigated technical requirements and challenges within this innovative sensing approach, detailing a design for crucial elements like sensing, networking, processing, and communication tailored for active mode research.

### **Network-constrained tracking of cyclists and pedestrians**

The increase in perception capabilities of connected mobile sensor platforms leads to an extensive surge of sensed features at various temporal and spatial scales. Harnessing available observations could enable to see how and where people move on sidewalks and cycle path. We propose a new method for advanced traffic applications, tracking an unknown and varying number of moving targets (e.g., pedestrians or cyclists) constrained by a road network, using mobile spatially distributed sensor platforms. In contrast to existing approaches that often consider short time horizons and are mostly designed for unconstrained motion in two- and three-dimensional space, where complexity tends to grow exponentially due to the number of hypotheses, we introduce the concept of network bound targets into the multi-target tracking problem, and hence to derive a network-constrained multi-hypotheses tracker (NC-MHT) utilizing the general position determined from the map and the position on the road segment. A simulation study shows that the method performs well in comparison to the standard MHT filter in free space. Results particularly highlight network-constraint effects for more efficient target predictions over extended periods of time, and in the simplification of the measurement association process, as compared to not utilizing a network structure.

### **A novel data source for advanced traffic management and control**

Traditionally, cyclist traffic data is collected via loop detectors, push buttons, smartphone apps, or through extra infrastructure like detector loops, cameras, and radars. However, mobile app usage is still low, and additional infrastructure is costly given the expenses related to deployment and maintenance. The main objective of this study was to show how augmenting the situational awareness of traffic signal controllers, using observations from moving sensor platforms

(*e.g.*, only considering position and speed), can enable prioritization of cyclists and reduce lost time within the control cycle in an effective way. We investigated the potential of using this new source of information, using a revised vehicle-actuated controller. This controller exploits CAV-generated observations of cyclists to optimize the control for cyclists. The results from a simulation study indicated that with a low CAV penetration rate, prioritizing cyclists by tracking reduced cyclist delays and stops, even with a small field of view. We point to future work, optimizing the control that balances the delays and stops of cyclists and cars.

### **Privacy challenges introduced by intelligent autonomous systems**

IAS are expected to coordinate data across time and space, introducing a mobile surveillance infrastructure into the environments that they function in. We foresee this newly introduced infrastructure to reconfigure the lived environment and see these developments as signaling an infrastructural build-out, the consequences of which will have privacy implications. We investigated the privacy implications that arise in traffic engineering scenarios within this sensing paradigm, particularly focusing on how people's movements are monitored. We used the term *non-user* to refer to someone who is in the environment of a CAV and captured by its sensing devices. The privacy of non-users (*e.g.*, pedestrians, cyclists outside the vehicle) in CAV settings should be a major source of concern, yet is a blind spot in the related literature. We identify and highlight key implications on location privacy for pedestrians and cyclists, marking the first effort to formalize privacy for active modes. The consistent distinction with other existing systems and the defined terminology can help researchers from different fields understand the intricacies of this complex socio-technical system and motivates for future research.

### **Conclusions and implications**

To the best of our knowledge, this thesis represents the first academic exploration into the potential utility and societal impact of exploiting pedestrian and cyclist data from moving sensor platforms, in the context of traffic research and practice.

At the technological level, the progress primarily driven by fields related to IAS demonstrate by implementation that available technology, effectively supports key aspects of the proposed sensing paradigm today. Thorough review of existing literature, however, reveals a notable absence of both theoretical frameworks and practical implementations for the sensing paradigm. The multitude of scenarios in which real-world data represents a potential asset for solving research questions necessitates a thoughtful mapping of objectives to sensing requirements and network characteristics, tailored to specific contexts and environments. Areas for future research include developing data fusion techniques, multi-object tracking, and scene understanding. Additionally, findings suggest potential for next-generation traffic surveillance applications, which presents opportunities for further research and development.

The sensing paradigm, however, is deeply intertwined with broader societal factors, including technological, social, and institutional considerations that collectively shape distinct usage patterns and raise important ethical concerns. This work displays that academic conversations surrounding the privacy of pedestrians and cyclists within IAS still seem to often overlook critical nuances, highlighting evident gaps in existing privacy and transport literature on this specific subject. This work also hints at broader risks extending beyond individual location privacy concerns. This work motivates to bridge the gap with the privacy engineering community to develop systems that balance utility and privacy, prioritizing approaches that embed ethical values by design.



# Samenvatting

Om veiligere, rechtvaardigere en milieuvriendelijkere vervoersopties te bevorderen, streven overheden ernaar om het aandeel verplaatsingen te voet en met de fiets te verhogen. Ondanks het erkende belang van onderzoek naar actieve vervoerswijzen, is er echter nog steeds een relatief gebrek aan empirische benaderingen, geavanceerde instrumenten en gegevens van hoge kwaliteit om de dynamiek van wanneer, waar en hoe mensen lopen en fietsen te begrijpen. Dit resulteerde in relatief beperkte wetenschappelijke vooruitgang en over het algemeen zeer beperkte implementatie van deze bevindingen in steden.

Dit proefschrift maakt gebruik van een nieuw detectieparadigma, aangedreven door de inzet van intelligente autonome systemen (IAS) in steden (zelfrijdende auto's, drones en andere instanties), wat resulteert in een flinke toename in het observeren van menselijke kenmerken — zoals hun bewegingspatronen, vormen en contextuele informatie — op grote temporele en ruimtelijke schaal. Belangrijk is dat dit proefschrift ook van toepassing is op andere bewegende sensorplatforms (MSP) die zijn uitgerust met geavanceerde sensormogelijkheden die niet afhankelijk zijn van deze informatie om zichzelf te kunnen besturen, zoals traditionele bussen, fietsen of zelfs mensen met bodycams. Wanneer ze autonoom zijn, zijn deze systemen echter specifiek ontworpen om spatiotemporele gegevens vast te leggen van alle objecten en verkeer in hun omgeving, inclusief voetgangers en fietsers, omdat ze deze informatie nodig hebben om zichzelf te kunnen besturen. De gegenereerde gegevens kunnen een waardevolle nieuwe bron van informatie vormen voor verkeersonderzoekers en beleidsmakers, maar hebben implicaties voor de privacy of een breder potentieel voor misbruik.

Met dit in gedachten is het doel van dit proefschrift *om het potentiële nut en de maatschappelijke impact te onderzoeken van het exploiteren van voetgangers- en fietsersgegevens van bewegende sensoren (IAS of MSP) in de context van verkeersonderzoek en -praktijk*. Het gebruik van de term *exploitatie* is tweeledig: 1) de ontwikkeling van methodologische kennis en hulpmiddelen om verkeersinformatie uit de gegevens te extraheren, en 2) onder welke voorwaarden (ijkpunten, beperkingen, risico's en schade) dit bijdraagt aan technologische, maatschappelijke en economische ontwikkeling.

Om dit doel te bereiken en een holistisch begrip te krijgen van dit opkomende paradigma, bouwt dit proefschrift voort op een interdisciplinaire onderzoeksbenadering, waarbij een brug wordt geslagen tussen transportonderzoek en andere relevante wetenschappelijke communities. Om een solide basis te leggen voor onze aannames en ontwerpkeuzes voor de toekomst, hebben we een theoretisch kader opgesteld door verschillende gegevensbronnen en kennis van experts samen te voegen. Deze conceptualisatie, die specifiek is toegesneden op onderzoek naar en de praktijk van de actieve modi, richt zich op de belangrijkste onderdelen van het systeem in sensor-, netwerk-, verwerkings- en communicatie-eenheden en stelde ons in staat om eisen en

functionele uitdagingen vast te stellen, en om onderzoek te schetsen dat moet worden uitgevoerd met de gegenereerde gegevens (hoofdstuk 2). Vervolgens hebben we onze analyse verfijnd door zowel de temporele als de ruimtelijke dimensies van de gegevens te onderzoeken, aangezien deze op verschillende niveaus van granulariteit en nauwkeurigheid konden worden gegenereerd. We hebben een methodologie ontwikkeld die consistent bijgewerkte locatie-informatie over de actieve modus oplevert voor het hele netwerk, waarmee we individuele trajecten over een wegennetwerk kunnen reconstrueren met behulp van mobiele, ruimtelijk gedistribueerde sensorplatforms (Hoofdstuk 3). Met behulp van deze aanpak en de inzichten die we tot nu toe hebben ontwikkeld, onderzoeken en demonstreren we hoe dit nieuwe type gegevens waardevolle inzichten kan opleveren en kan helpen bij het optimaliseren van verkeersregel- en managementstrategieën. We definieerden een verkeersregelscenario en gebruikten modellering en simulatie om scenariospecifieke gegevens op schaal te genereren (hoofdstuk 4). Door het gebrek aan adequate gegevens uit de echte wereld, werd waar nodig gebruik gemaakt van simulatie om de onderzoeksvragen 2 en 3 te beantwoorden. Bij het verzamelen van gegevens over individuen is er echter een inherente wisselwerking tussen nut en privacy. De bredere socio-technische impact van het detectieparadigma wordt onderzocht door het te vergelijken met huidige systemen, waarbij de nadruk ligt op de gevolgen voor de locatieprivacy van niet-gebruikers (voetgangers en fietsers) (hoofdstuk 5). De volgende paragrafen beschrijven de bevindingen van dit onderzoek met betrekking tot deze vier onderwerpen.

### **Naar nieuwe systeemontwerpen voor actieve mobiele detectie**

We presenteren een theoretisch ontwerp voor een nieuw mobiel sensorsysteem (AMSense) dat voertuigen gebruikt als mobiele sensorknooppunten in een netwerk om spatiotemporele eigenschappen van voetgangers en fietsers in stedelijke omgevingen vast te leggen. In deze dynamische multi-sensor benadering worden real-time gegevens, algoritmen en modellen samengevoegd om de aanwezigheid, posities en bewegingen van actieve modi in te schatten met informatie die wordt gegenereerd door een vloot van mobiele sensorplatforms. Het is een schaalbaar systeem dat antwoorden geeft op spatiotemporele resolutie, penetratie en dynamische netwerkomstandigheden. Het bewegende platform kan dus de robuustheid van de observatie vergroten door gegevens van meerdere bronnen te gebruiken. In dit hoofdstuk motiveren we de noodzaak en tonen we het potentieel van een dergelijk detectieparadigma, dat een groot aantal nieuwe onderzoeks- en toepassingsontwikkelingen ondersteunt. We hebben de technische vereisten en uitdagingen binnen deze innovatieve detectiebenadering onderzocht en een ontwerp voor cruciale elementen zoals detectie, netwerken, verwerking en communicatie op maat gemaakt voor onderzoek naar de actieve modi.

### **Netwerkbeperkt volgen van fietsers en voetgangers**

De toename in waarnemingsmogelijkheden van verbonden mobiele sensorplatforms leidt tot een grote hoeveelheid waargenomen kenmerken op verschillende tijd- en ruimteschalen. Door gebruik te maken van de beschikbare waarnemingen kan men zien hoe en waar mensen zich bewegen op trottoirs en fietspaden. Wij stellen een nieuwe methode voor geavanceerde verkeerstoeepassingen voor, waarbij een onbekend en variërend aantal bewegende doelen (bijv. voetgangers of fietsers), begrensd door een wegennetwerk, wordt gevolgd met behulp van mobiele, ruimtelijk gedistribueerde sensorplatforms. In tegenstelling tot bestaande benaderingen die vaak rekening houden met een korte tijdshorizon en meestal ontworpen zijn voor onbepaalde beweging in twee- en driedimensionale ruimte, waarbij de complexiteit exponentieel toeneemt door het aantal hypothesen, introduceren we het concept van netwerkgebonden doelen in het

multitarget-trackingprobleem, en leiden we zo een netwerkbepaalde multi-hypothesetracker (NC-MHT) af die gebruikmaakt van de algemene positie die bepaald wordt uit de kaart en de positie op het wegsegment. Een simulatiestudie toont aan dat de methode goed presteert in vergelijking met het standaard MHT-filter in de vrije ruimte. De resultaten benadrukken met name de netwerkbepaalde effecten voor efficiëntere doelvoorspellingen over langere perioden en in de vereenvoudiging van het meetassociatieproces, vergeleken met het niet gebruiken van een netwerkstructuur.

### **Een nieuwe gegevensbron voor geavanceerd verkeersbeheer en -regeling**

Traditioneel worden verkeersgegevens van fietsers verzameld via lusedetectoren, drukknoppen, smartphone apps of via extra infrastructuur zoals detectielussen, camera's en radars. Het gebruik van mobiele apps is echter nog steeds laag en extra infrastructuur is kostbaar gezien de kosten voor installatie en onderhoud. Het hoofddoel van dit onderzoek was om aan te tonen hoe het vergroten van het situationeel bewustzijn van verkeersregelaars, met behulp van observaties van bewegende sensorplatforms (alleen rekening houdend met positie en snelheid), het mogelijk kan maken om fietsers voorrang te geven en verloren tijd binnen de regelcyclus op een effectieve manier te verminderen. We onderzochten het potentieel van het gebruik van deze nieuwe informatiebron met behulp van een aangepaste voertuiggestuurde regelaar. Deze regelaar maakt gebruik van CAV-waarnemingen van fietsers om de regeling voor fietsers te optimaliseren. De resultaten van een simulatiestudie toonden aan dat bij een lage penetratiegraad van CAV, het prioriteren van fietsers, door het volgen van fietsers, de vertragingen en stops van fietsers verminderde, zelfs met een klein gezichtsveld. We wijzen op toekomstig werk, het optimaliseren van de regeling die de vertragingen en stops van fietsers en auto's in balans brengt.

### **Privacy-uitdagingen geïntroduceerd door intelligente autonome systemen**

Van IAS wordt verwacht dat ze gegevens in tijd en ruimte coördineren, waardoor ze een mobiele bewakingsinfrastructuur introduceren in de omgevingen waarin ze functioneren. We verwachten dat deze nieuw geïntroduceerde infrastructuur de leefomgeving zal herconfigureren en zien deze ontwikkelingen als het signaal van een infrastructurele opbouw, waarvan de gevolgen impact zullen hebben op de privacy. We onderzochten de privacyimplicaties die zich voordoen in verkeerskundige scenario's binnen dit detectieparadigma, waarbij we ons met name richtten op de manier waarop de bewegingen van mensen worden gemonitord. We gebruikten de term *niet-gebruiker* om te verwijzen naar iemand die zich in de omgeving van een CAV bevindt en opgevangen wordt door de sensoren. De privacy van niet-gebruikers (voetgangers, fietsers buiten het voertuig) in CAV omgevingen zou een grote bron van zorg moeten zijn, maar is een blinde vlek in de gerelateerde literatuur. We identificeren en benadrukken de belangrijkste implicaties voor de locatieprivacy van voetgangers en fietsers, als een eerste poging om privacy voor actieve modi te formaliseren. Het consistente onderscheid met andere bestaande systemen en de gedefinieerde terminologie kan onderzoekers uit verschillende vakgebieden helpen de fijne kneepjes van dit complexe socio-technische systeem te begrijpen en motiveert toekomstig onderzoek.

### **Conclusies en implicaties**

Voor zover wij weten, vertegenwoordigt dit proefschrift de eerste academische verkenning naar het potentiële nut en de maatschappelijke impact van het gebruik van gegevens van voetgangers en fietsers van bewegende sensorplatforms in de context van verkeersonderzoek en -praktijk.

Op technologisch niveau toont de vooruitgang die voornamelijk wordt gedreven door vakgebieden die gerelateerd zijn aan IAS, door implementatie aan dat de beschikbare technologie vandaag de dag effectief belangrijke aspecten van het voorgestelde detectieparadigma ondersteunt. Een grondige studie van de bestaande literatuur onthult echter een opvallende afwezigheid van zowel theoretische kaders als praktische implementaties voor het detectieparadigma. De vele scenario's waarin gegevens uit de echte wereld een potentiële bron zijn voor het oplossen van onderzoeksvragen, vereisen dat doelstellingen naar detectievereisten en netwerkenmerken doordacht in kaart worden gebracht en op maat gemaakt voor specifieke contexten en omgevingen. Gebieden voor toekomstig onderzoek zijn onder andere de ontwikkeling van datafusietechnieken, het volgen van meerdere objecten en het begrijpen van specifieke situaties. Daarnaast suggereren de bevindingen een potentieel voor de volgende generatie verkeersbewakingstoepassingen, wat mogelijkheden biedt voor verder onderzoek en ontwikkeling.

Het detectieparadigma is echter sterk verweven met bredere maatschappelijke factoren, waaronder technologische, sociale en institutionele overwegingen die samen vorm geven aan verschillende gebruikspatronen en die belangrijke ethische vragen oproepen. Dit werk laat zien dat academische gesprekken over de privacy van voetgangers en fietsers binnen het IAS nog steeds vaak kritieke nuances over het hoofd lijken te zien, en benadrukt duidelijke hiaten in de bestaande privacy- en transportliteratuur over dit specifieke onderwerp. Dit werk wijst ook op bredere risico's die verder gaan dan de privacy van de individuele locatie. Dit werk motiveert om de kloof met de privacy engineering gemeenschap te overbruggen om systemen te ontwikkelen die nut en privacy in evenwicht brengen, waarbij prioriteit wordt gegeven aan benaderingen die ethische waarden verankeren in het ontwerp.



# About the author

Alphonse Vial is a researcher exploring the many applications of emerging data-driven technologies, with a strong focus on their often-overlooked impact on individuals, workflows, and broader societal dynamics. In order to adequately explore these challenges, a significant part of his research is interdisciplinary.

He has a background in Mechatronics and Technology Management, and conducted his PhD in the department of Transport and Planning at Delft University of Technology. His research interests began with data-driven urban transportation systems and machine learning. Over time, they evolved to include sensor fusion, privacy, and, more recently, neurosymbolic and generative approaches, as well as the broader social implications of AI and algorithmic systems. Along the way, he has been fortunate to collaborate with and learn from researchers at institutions such as TU Delft, Linköping University, École Polytechnique Fédérale de Lausanne, Massachusetts Institute of Technology, Fraunhofer Institutes, Eidgenössische Technische Hochschule Zürich, and the University of Amsterdam, experiences that have helped shape and refine his research interests.

Alphonse has authored academic papers, has contributed to major national and European grant proposals, and has served as a peer reviewer for leading journals and conferences in his field. He has faced his share of rejections, extensive revisions, and the occasional harsh reviewer comment — but he still genuinely enjoys being part of the scientific community, working across disciplines, and staying curious.



# Publications

## Journal papers

- **Vial, A.**, Daamen, W., Ding, A.Y., van Arem, B., and Hoogendoorn, S.P. (2020), AM-Sense: How Mobile Sensing Platforms Capture Pedestrian/Cyclist Spatiotemporal Properties in Cities. *IEEE Intelligent Transportation Systems Magazine*, 14(1), pp. 29-43.
- **Vial, A.**, Hendebay, G., Daamen, W., van Arem, B., and Hoogendoorn, S.P. (2022). Framework for Network-Constrained Tracking of Cyclists and Pedestrians. *IEEE Transactions on Intelligent Transportation Systems*, 24(3), pp. 3282-3296.
- **Vial, A.**, Salomons, A.M., Daamen, W., van Arem, B., Hoogendoorn-Lanser, S., and Hoogendoorn, S.P. (2023). Prioritizing cyclists at signalized intersections using observations from moving sensors. *Transportation Research Record: Journal of the Transportation Research Board*, 2677(12), 29-43.

## Conference contributions

- **Vial A.**, Schmidt, A. (2019). Quantifying the potential of electrification with large-scale vehicle trajectory data. *mobil.TUM 2018, International Scientific Conference on Mobility and Transport* (Poster award).
- Zomer, L.B., **Vial, A.**, Reggiani, G., Ton, D., Wierbos, M.J., Gong, V., Schneider, F., Feng, Y., Sparnaaij, M., Gavriilidou, A., van Oijen, T., Duives, D.C., Yuan, Y., Cats, O., Knoop, V.L., Daamen, W. and Hoogendoorn, S.P. (2019). The Impact of Cycling Research: Connecting Science and Practice. Presented at Cycling Research Board, October 2019, Delft, The Netherlands.

# TRAIL Thesis Series

The following list contains the most recent dissertations in the TRAIL Thesis Series. For a complete overview of more than 400 titles, see the website: [www.rsTRAIL.nl](http://www.rsTRAIL.nl).

The TRAIL Thesis Series is a series of the Netherlands TRAIL Research School on transport, infrastructure and logistics.

Vial, A., *Eyes in Motion: A new traffic sensing paradigm for pedestrians and cyclists*, T2025/11, May 2025, TRAIL Thesis Series, the Netherlands

Chen, Q., *Towards Mechanical Intelligence in Soft Robotics: Model-based design of mechanically intelligent structures*, T2025/10, April 2025, TRAIL Thesis Series, the Netherlands

Eftekhari, Z., *Exploring the Spatial and Temporal Patterns in Travel Demand: A data-driven approach*, T2025/9, June 2025, TRAIL Thesis Series, the Netherlands

Reddy, N., *Human Driving Behavior when Interacting with Automated Vehicles and the Implications on Traffic Efficiency*, T2025/8, May 2025, TRAIL Thesis Series, the Netherlands

Durand, A., *Lost in Digitalisation? Navigating public transport in the digital era*, T2025/7, May 2025, TRAIL Thesis Series, the Netherlands

Dong, Y., *Safe, Efficient, and Socially Compliant Automated Driving in Mixed Traffic: Sensing, Anomaly Detection, Planning and Control*, T2025/6, May 2025, TRAIL Thesis Series, the Netherlands

Droffelaar, I.S. van, *Simulation-optimization for Fugitive Interception*, T2025/5, May 2025, TRAIL Thesis Series, the Netherlands

Fan, Q., *Fleet Management Optimisation for Ride-hailing Services: from mixed traffic to fully automated environments*, T2025/4, April 2025, TRAIL Thesis Series, the Netherlands

Hagen, L. van der, *Machine Learning for Time Slot Management in Grocery Delivery*, T2025/3, March 2025, TRAIL Thesis Series, the Netherlands

Schilt, I.M. van, *Reconstructing Illicit Supply Chains with Sparse Data: a simulation approach*, T2025/2, January 2025, TRAIL Thesis Series, the Netherlands

Ruijter, A.J.F. de, *Two-Sided Dynamics in Ridesourcing Markets*, T2025/1, January 2025, TRAIL Thesis Series, the Netherlands

Fang, P., *Development of an Effective Modelling Method for the Local Mechanical Analysis of Submarine Power Cables*, T2024/17, December 2024, TRAIL Thesis Series, the Netherlands

Zattoni Scroccaro, P., *Inverse Optimization Theory and Applications to Routing Problems*, T2024/16, October 2024, TRAIL Thesis Series, the Netherlands

Kapousizis, G., *Smart Connected Bicycles: User acceptance and experience, willingness to pay and road safety implications*, T2024/15, November 2024, TRAIL Thesis Series, the Netherlands

Lyu, X., *Collaboration for Resilient and Decarbonized Maritime and Port Operations*, T2024/14, November 2024, TRAIL Thesis Series, the Netherlands

Nicolet, A., *Choice-Driven Methods for Decision-Making in Intermodal Transport: Behavioral heterogeneity and supply-demand interactions*, T2024/13, November 2024, TRAIL Thesis Series, the Netherlands

Kougiatsos, N., *Safe and Resilient Control for Marine Power and Propulsion Plants*, T2024/12, November 2024, TRAIL Thesis Series, the Netherlands

Uijtdewilligen, T., *Road Safety of Cyclists in Dutch Cities*, T2024/11, November 2024, TRAIL Thesis Series, the Netherlands

Liu, X., *Distributed and Learning-based Model Predictive Control for Urban Rail Transit Networks*, T2024/10, October 2024, TRAIL Thesis Series, the Netherlands

Clercq, G. K. de, *On the Mobility Effects of Future Transport Modes*, T2024/9, October 2024, TRAIL Thesis Series, the Netherlands

Dreischerf, A.J., *From Caveats to Catalyst: Accelerating urban freight transport sustainability through public initiatives*, T2024/8, September 2024, TRAIL Thesis Series, the Netherlands

Zohoori, B., *Model-based Risk Analysis of Supply Chains for Supporting Resilience*, T2024/7, October 2024, TRAIL Thesis Series, the Netherlands



TRAIL

## **Summary**

The deployment of moving sensor platforms (e.g., self-driving cars, drones, and other instances) with advanced sensing is rapidly increasing the capture of human features at unprecedented temporal and spatial scales, especially in cities. This thesis advances knowledge on extracting information from this novel data source for traffic research and practice, while highlighting implications for privacy and beyond. Findings provide insights for traffic control and management, policy development, and anyone involved in responsible urban innovation.

## **About the Author**

Alphonse Vial conducted his PhD at TU Delft and has a background in Technology Management and in Mechatronics. His research explores the opportunities and challenges of emerging data-driven technologies through an interdisciplinary lens.

**TRAIL Research School** ISBN 978-90-5584-366-4



ERASMUS UNIVERSITY ROTTERDAM

Radboud University



rijksuniversiteit  
 groningen



TU Delft  
 Delft University of  
 Technology

UNIVERSITY OF TWENTE.



Technische Universiteit  
 Eindhoven  
 University of Technology