

Titre: Driver Behaviour and Road Safety Analysis Using Computer Vision
Title: and Applications in Roundabout Safety

Auteur: Paul St-Aubin
Author:

Date: 2016

Type: Mémoire ou thèse / Dissertation or Thesis

Référence: St-Aubin, P. (2016). Driver Behaviour and Road Safety Analysis Using Computer
Citation: Vision and Applications in Roundabout Safety [Thèse de doctorat, École
Polytechnique de Montréal]. PolyPublie. <https://publications.polymtl.ca/2272/>

 **Document en libre accès dans PolyPublie**
Open Access document in PolyPublie

URL de PolyPublie: <https://publications.polymtl.ca/2272/>
PolyPublie URL:

**Directeurs de
recherche:** Nicolas Saunier, & Luis Miranda-Moreno
Advisors:

Programme: Génie civil
Program:

UNIVERSITÉ DE MONTRÉAL

DRIVER BEHAVIOUR AND ROAD SAFETY ANALYSIS USING COMPUTER VISION
AND APPLICATIONS IN ROUNDABOUT SAFETY

PAUL ST-AUBIN
DÉPARTEMENT DES GÉNIES CIVIL, GÉOLOGIQUE ET DES MINES
ÉCOLE POLYTECHNIQUE DE MONTRÉAL

THÈSE PRÉSENTÉE EN VUE DE L'OBTENTION
DU DIPLÔME DE PHILOSOPHIÆ DOCTOR
(GÉNIE CIVIL)
AOÛT 2016

UNIVERSITÉ DE MONTRÉAL

ÉCOLE POLYTECHNIQUE DE MONTRÉAL

Cette thèse intitulée :

DRIVER BEHAVIOUR AND ROAD SAFETY ANALYSIS USING COMPUTER VISION
AND APPLICATIONS IN ROUNDABOUT SAFETY

présentée par: ST-AUBIN Paul

en vue de l'obtention du diplôme de: Philosophiæ Doctor

a été dûment acceptée par le jury d'examen constitué de:

M. TRÉPANIÉ Martin, Ph. D., président

M. SAUNIER Nicolas, Ph. D., membre et directeur de recherche

M. MIRANDA-MORENO Luis, Ph. D., membre et codirecteur de recherche

M. DANIELS Stijn, Ph. D., membre

Mme PARK Seri, Ph. D., membre

DEDICATION

*To my teachers and mentors,
family and friends. . .*

ACKNOWLEDGEMENTS

This research project was initiated as part of the Safety Study of Québec Roundabouts project (*Sécurité des carrefours giratoires au Québec*, proposal number 2012-SO-163493, 2012-2014) funded by the Fonds de recherche du Québec – Nature et technologies (FRQNT), and the Fonds de recherche du Québec – Santé (FRQS), and funded and overseen by the Ministère des Transports du Québec (MTQ), as part of their research program on road safety. The author wishes to acknowledge the help of participating municipalities in the province of Québec and in Sweden, and the MTQ, for technical and logistical support during data collection. The Safety Study of Québec Roundabouts projects was headed by Nicolas Saunier from Polytechnique Montréal, and included participating researchers at other universities including Luis F. Miranda-Moreno from McGill University, Zachary Patterson from Concordia University, Karim Ismail from Carleton University, as well as participating students at each university.

This research project was furthered and extended thanks to a Natural Sciences and Engineering Research Council of Canada (NSERC) doctoral research grant (PGS.D) (2014-2016), a FRQNT doctoral research grant (2014-2015), and a complementary doctoral scholarship from the Réseau de recherche en sécurité routière (RRSR) (2013).

I am indebted to my supervisors, Nicolas Saunier from Polytechnique Montréal, and Luis F. Miranda-Moreno from McGill University, who also supervised my Master's Thesis. Their guidance and support was a source of both inspiration and determination. I wish to thank my mother, D.M. St-Aubin, for proofreading, driving my pursuit of higher education, and everything else.

An international doctoral exchange was made possible thanks to Aliaksei Lareshyn at Lund University, Sweden, and was funded by a RRSR student exchange scholarship (2015).

Many graduate students were instrumental in helping conduct data collection across the province of Québec: Shaun Burns, Arthur Dolmajian, Matin Nabavi, and Joshua Stipanovic. Philip Morse and Laurent Gauthier contributed some code to *tvaLib*.

RÉSUMÉ

L'un des principaux défis provenant de l'analyse traditionnelle de la sécurité routière basée sur les données historiques d'accidents est le besoin d'observer de véritables collisions entre les usagers de la route. Non seulement indésirables, ces collisions sont difficiles à observer. À cet effet, les méthodes d'analyse substitutive de la sécurité routière gagnent du terrain dans le milieu de la recherche en tant qu'alternative proactive à l'observation de ces accidents de la route : cette approche promet de modéliser indirectement les collisions par l'intermédiaire de précurseurs de collisions pouvant être retrouvés dans des données de circulation ordinaire : situations de trajectoire de collision, quasi-accidents, conflits de circulation, etc. En sus, l'analyse substitutive de la sécurité donne également les chercheurs un aperçu des mécanismes de collision, ce qui permettrait de mieux comprendre les facteurs favorisant les accidents. Cependant, un grand nombre de définitions de ces situations précurseuses de collision, ainsi que des problèmes de cohérence et de subjectivité des méthodes de collecte de données ont entravé l'adoption des méthodes d'analyse substitutive de la sécurité routière.

Dans cette thèse, ces problèmes de cohérence de données et de définitions de précurseurs de collision sont adressés de manière systématique. À ces fins, un système automatisé de collecte de données de circulation de haute résolution (jusqu'à 30 observations par seconde par usager de la route) à grande taille (près de 20,000 véh-km) à partir d'extraits vidéo est développé et validé. Ensuite, un cadre d'analyse substitutive de la sécurité systématique est mis au point. Enfin, ce cadre est appliqué à un cas d'étude de 51 zones d'entrecroisement de carrefours giratoires à travers le Québec et à Lund, en Suède. Cette comparaison internationale du comportement des usagers de la route et de la sécurité constitue une première du genre à ce niveau de détail.

Un système automatisé de collecte de données de circulation à partir d'extraits vidéo est construit grâce aux améliorations récentes dans le domaine de la vision par ordinateur, en particulier en ce qui concerne les algorithmes de suivi de mouvement et de correction d'image. Ceci permet d'obtenir un ensemble de données systématiquement mesurées qui sont à la fois détaillés et à grande portée. Ces trajectoires sont validées en utilisant un ensemble de techniques d'optimisation de suivi des trajectoires et de correction des erreurs. Les précurseurs de collision classiques utilisés pour l'analyse substitutive de sécurité, tel que la vitesse, le créneau au céder-le-passage, et le temps-à-la-collision sont ensuite adaptés et calculés pour les données recueillies. Le temps-à-la-collision, qui dépend de modèles de prédiction des trajectoires et de collisions potentielles, est étudié à l'aide de modèles classiques de prédic-

tion du mouvement qui ne tienne pas compte du contexte (géométrie des aménagements, circulation). Puisque ces modèles de prédiction de mouvement ne se prêtent pas bien à la modélisation de mouvement en situations non-linéaires, telles que l'on peut trouver dans les zones d'entrecroisement des carrefours giratoires, un nouveau modèle de prédiction de mouvement des usagers de la route basé sur les patrons de mouvement est développé. Ce modèle naturaliste demeure relativement peu coûteux à calculer et simple à mettre en œuvre.

Pour la sélection des carrefours giratoires, un échantillon de sites à travers la province du Québec est sélectionné pour représenter les différentes façons dont les zones d'entrecroisement sont conçues et la variété d'occupation du sol aux alentours des carrefours giratoires. En ce qui concerne l'étude internationale, les sites sont soigneusement choisis pour que les facteurs de conception géométrique, d'occupation du sol et de conditions de circulation soient aussi similaires que possibles en Suède et au Québec.

La précision du suivi des trajectoires à partir des algorithmes de vision par ordinateur est mesurée selon le *Multiple Object Tracking Accuracy*. L'optimisation générale des paramètres de suivi résulte en une précision de suivi mesurée entre 80 % et 85 %, avant l'application de mesures de corrections d'erreur. Le nouveau modèle de prédiction de mouvement à base de patron de mouvement possède plusieurs propriétés intrinsèques souhaitables et se révèle comme étant réalisable, et une agrégation du 15^e centile des mesures temps-à-la-collision au cours des paires d'usagers de la route demeure nécessaire.

En ce qui concerne les carrefours giratoires, on démontre qu'un certain nombre de paramètres de conception géométrique et d'utilisation du sol ont une corrélation avec les différentes mesures substitutive de sécurité. Les mesures substitutives de sécurité (chacune représentant un aspect différent de la sécurité routière) sont jugées être indépendantes les unes les autres, et la plupart des facteurs explicatifs sont associés à différents aspects de la sécurité routière. De plus, l'analyse internationale de mesures substitutives de sécurité montre des résultats cohérents avec le bilan routier de ces deux régions : après contrôle de l'exposition, la Suède enregistre enregistre moitié moins d'accidents et de décès sur la route que le Québec, et cela se reflète dans l'observation en Suède de mesures de vitesses systématiquement inférieures et de mesures de créneau et de temps-à-la-collision systématiquement supérieures.

ABSTRACT

One of the main challenges of traditional road safety analysis based on historical accident records is its dependence on the occurrence and subsequent observation of real traffic collisions. Traffic collisions are not only undesirable, they are difficult to observe. Surrogate safety analysis is gaining traction in the research community as a proactive alternative to observing historical accident records. With this approach, collisions are instead predicted indirectly via precursors to collisions found in everyday traffic scenarios: collision-courses, near-misses, traffic conflicts, etc. Furthermore, the scope of surrogate safety analysis provides insight into collision mechanisms, allowing for better investigative procedures. However, the wide range of collision precursor definitions, and issues with inconsistent or sometimes subjective data collection methods have hampered surrogate safety analysis adoption.

In this thesis, the issues of data quality and systematic collision precursor definitions are addressed. To this end, an automated video-based traffic data collection system capable of extracting large quantities (nearly 20,000 veh-km) of high-resolution road user trajectories (up to 30 observations per second per road user) is first built and validated. Then, a precisely defined surrogate safety framework which makes use of the trajectory data is devised. Finally, this framework is showcased in a case study of 51 roundabout merging zones across Québec, Canada, and Lund, Sweden. The international comparison of road user behaviour and safety is a first of its kind at this level of detail.

An automated video-based traffic data collection system is built using recent improvements in camera technology, computer vision, and image correction algorithms. This provides a set of systematically measured data which do not compromise study scale and level of detail. These trajectories are validated using a battery of new tracking accuracy optimization and error filtering techniques. Classic surrogate safety measures of speed, post-encroachment time, and time-to-collision are then adapted from the literature to be measured from the newly obtained, detailed trajectory data. Time-to-collision, which depends on collision course modelling, is investigated using classic, context-free motion prediction models. However, these context-free motion prediction models do not lend themselves well to modelling complex environments as can be found in roundabout merging zones. As such, a new motion-pattern-based motion prediction model is developed which is capable of modelling naturalistic road user behaviour, yet remains relatively inexpensive to calculate and simple to implement and expand.

A sample of sites across the province of Québec is selected to represent a large variety of

roundabout designs and land uses. For the Swedish sites, site selection is carefully performed to control for as many geometric designs, land uses, and traffic parameters as possible both in Sweden and Québec.

Trajectory data are validated using Multiple Object Tracking Accuracy metrics. General-purpose parameter optimization results in tracking accuracies in the low 80 % range, before error correction is applied. The new motion-pattern-based motion prediction model, possessing several intrinsic properties making it more desirable, is shown to be feasible, and an aggregation of time-to-collision measures over pairs of road users is deemed necessary, opting for the 15th centile measure.

As for roundabouts, a number of geometric and land use parameters are shown to be correlated with surrogate safety measures. The surrogate safety measures themselves, each representing a different aspect of road safety, are found to be largely independent of each other, and that most explanatory factors seem to affect different aspects of road safety. Results of the comparison of surrogate safety measures in Québec and Sweden are consistent with historical accident records: after controlling for exposure, Sweden records half as many accidents and traffic-collision fatalities as Québec, and this is reflected in systematically lower speeds and higher post-encroachment time and time-to-collision measures recorded in Sweden.

TABLE OF CONTENTS

DEDICATION	iii
ACKNOWLEDGEMENTS	iv
RÉSUMÉ	v
ABSTRACT	vii
TABLE OF CONTENTS	ix
LIST OF TABLES	xiii
LIST OF FIGURES	xv
CHAPTER LIST OF SYMBOLS AND ABBREVIATIONS	xviii
CHAPTER GLOSSARY	xx
LIST OF APPENDICES	xxviii
CHAPTER 1 INTRODUCTION	1
1.1 Research Objectives	6
1.1.1 Contributions	7
1.1.2 Limitations and Scope	8
1.2 Outline	9
CHAPTER 2 LITERATURE REVIEW	11
2.1 Video-Based Traffic Data Collection Systems	11
2.1.1 Overview of Automated Computer Vision Systems	12
2.2 Safety Framework	15
2.2.1 Collision Probability	16
2.2.2 Collision Severity	18
2.3 Surrogate Safety Methods	19
2.3.1 History and Purpose	19
2.3.2 Modern Use of Surrogate Safety Methods	20
2.3.3 Collision-Course Motion Prediction	23

2.4	Roundabouts	26
2.4.1	Roundabout Safety Studies	27
CHAPTER 3 A NOVEL FRAMEWORK FOR LARGE-SCALE HIGH-RESOLUTION		
	TRAFFIC DATA COLLECTION USING COMPUTER VISION	32
3.1	System Overview	33
3.2	Video Data Collection	36
3.2.1	Equipment	37
3.2.2	Source Video Data Treatment	39
3.2.3	Video Data Storage and Indexing	42
3.3	Vehicle Tracking	42
3.3.1	Feature-Based Tracking	42
3.3.2	Data Structures	44
3.3.3	Validation	51
3.4	Semi-Supervised Model for Automated Traffic Analysis	52
3.4.1	Video Metadata	55
3.4.2	Alignments	56
3.4.3	High-Level Interpretation	57
3.4.4	Technical Improvements for Large-Scale Data Processing	57
3.4.5	Tracking Performance	59
3.4.6	Quality Control and Error Correction	60
3.4.7	Playback	66
3.5	Summary	68
CHAPTER 4 SURROGATE SAFETY FRAMEWORK FOR ROAD SAFETY ANAL-		
	YSIS	70
4.1	Introduction to Surrogate Safety	70
4.1.1	User Pairs and Interaction Instants	72
4.2	Directly Observable Surrogate Safety Measures	75
4.3	Collision-Course-Based Measures	76
4.3.1	Context-Free Motion Prediction	82
4.3.2	Prediction Based on Motion Pattern	83
4.3.3	Motion Prediction Methods Summary	89
4.4	Classification of User Pairs	91
4.5	Surrogate Safety Measure Aggregation Methods	93
4.6	Evaluating Road Safety Using Surrogate Safety Measures	95
4.6.1	Choice of Safety Evaluation	96

4.6.2	Comparison and Choice of TTC Measures	98
4.6.3	Summary of Surrogate Safety Measures	112
4.7	Summary	113
CHAPTER 5 CASE STUDY OF QUÉBEC ROUNDABOUTS		116
5.1	Roundabout Design Elements	116
5.1.1	Merging Zone as a Unit of Analysis	117
5.1.2	High-Level Interpretation Module: Roundabouts	120
5.2	Québec Roundabout Data Inventory	125
5.2.1	Site Selection	126
5.2.2	Data Inventory	131
5.2.3	Potential Contributing Factors	131
5.3	Overview and Exploratory Analysis	136
5.3.1	Flows	137
5.3.2	Speed	139
5.3.3	Yielding Post-Encroachment Time	142
5.3.4	Time-to-Collision	144
5.3.5	Analysis of Correlation of Factors	150
5.4	Regression Analysis	154
5.4.1	Traffic Parameters	156
5.4.2	Speed	158
5.4.3	Yielding Post-Encroachment Time	160
5.4.4	Time-to-Collision	161
5.5	Summary	164
CHAPTER 6 INTERNATIONAL COMPARISON BETWEEN QUÉBEC AND SWEDISH ROAD USERS AT ROUNDABOUTS		167
6.1	Rationale	168
6.2	Swedish Roundabout Data Inventory	170
6.2.1	Site Selection	170
6.2.2	Data Inventory	172
6.2.3	Potential Contributing Factors	179
6.2.4	Analysis of Correlation of Factors	180
6.3	Overview and Exploratory Analysis	181
6.3.1	Speed	181
6.3.2	Yielding Post-Encroachment Time	182
6.3.3	Time-to-Collision	183

6.4	Regression Analysis	184
6.4.1	Speed	184
6.4.2	Yielding Post-Encroachment Time	185
6.4.3	Time-to-Collision	185
6.5	Summary	186
CHAPTER 7 CONCLUSION		189
7.1	Summary Conclusions	189
7.1.1	Contributions to Applications of Computer Vision for Traffic Analysis	189
7.1.2	Contributions to Surrogate Safety Measures	190
7.1.3	Contributions to Roundabout Safety	192
7.2	Limitations	193
7.2.1	Applications of Computer Vision for Traffic Analysis	193
7.2.2	Surrogate Safety Measures	194
7.2.3	Roundabout Safety	194
7.3	Future Work	195
7.3.1	Applications of Computer Vision to Traffic Analysis	195
7.3.2	Surrogate Safety Measures	196
7.4	Final Perspective	198
BIBLIOGRAPHY		199
APPENDICES		216

LIST OF TABLES

Table 3.1	Tracking Performance	60
Table 3.2	Optimized General Summer Feature Tracker Parameters	61
Table 3.3	Summary of Default Filtering Parameters	66
Table 4.1	Summary of Prediction Methods Used	89
Table 4.2	Summary of Collision Course and Evasion Parameters	90
Table 4.3	Summary of Aggregation Methods (TTC)	95
Table 4.4	Spearman’s Ranking Correlation for Serious Event Comparison . . .	100
Table 4.5	Spearman’s Ranking Correlation for Safety Continuum Comparison .	112
Table 4.6	Safety Indicators	113
Table 4.7	Measures of Exposure or Related to Exposure	114
Table 5.1	Spearman’s Correlation of Ranking Based on Mean Site yPET below Threshold Value	123
Table 5.2	Indicators Characterizing Demand and Traffic Ratios for HLI at Round- abouts	125
Table 5.3	Québec Roundabout Data Inventory	131
Table 5.4	Geometric Factors	136
Table 5.5	K-means Cluster Profiles	150
Table 5.6	Correlation table for the remaining traffic, geometric, and environmen- tal factors (a)	152
Table 5.7	Correlation table for the remaining traffic, geometric, and environmen- tal factors (b)	153
Table 5.8	Spearman’s Correlation of Merging Zone-Aggregated Surrogate Safety Measure Indicators	156
Table 5.9	Linear Regression Models for Exposure	158
Table 5.10	Linear Regression Model for Mean Road User Speed	159
Table 5.11	Linear Regression Model for Median Lag yPET below 5 s	160
Table 5.12	Random-Effect Regression Models for Motion Pattern-Based 15 th Cen- tile Time-to-Collision (SCC)	163
Table 5.13	Linear Regression Models for Motion Pattern-Based Time-to-Collision Events Below 1.5 s Per Hour (Weighted SEC)	164
Table 5.14	Regression Model Summary	166
Table 6.1	Comparison of Macroscopic Historical Accident Statistics	168
Table 6.2	Sweden Roundabout Data Inventory	173

Table 6.3	Site Characteristic Inventory: Sweden	177
Table 6.4	Site Characteristic Inventory: Québec	178
Table 6.5	Additional Factors	179
Table 6.6	Accident Data	179
Table 6.7	Pearson’s linear correlation table for the traffic, geometric, and environmental factors	181
Table 6.8	Linear Regression Models for Mean Speed and Median Lag Yielding Post-Encroachment Time	185
Table 6.9	Random Effects Regression Models for Motion Pattern-Based 15 th Centile Time-to-Collision (SCC)	186
Table C.1	Stata Regression Commands and PostEstimation Statistics	224
Table D.1	Site Selection Scale Parameters	225

LIST OF FIGURES

Figure 1.1	Outline of the overall proposed methodology and results	10
Figure 2.1	Conflict point comparison between a regular intersection and a roundabout	28
Figure 3.1	Overview of the data collection pipeline embedded in a typical traffic engineering study	35
Figure 3.2	Example video data collection unit installation.	39
Figure 3.3	Side-by-side comparison of lens distortion correction	41
Figure 3.4	Example feature tracking of two vehicles	43
Figure 3.5	Close-up of image space for each RGB channel	45
Figure 3.6	Illustration of a homography projection	47
Figure 3.7	Comparison of Cartesian coordinates and curvilinear coordinates in world space	49
Figure 3.8	Moving object data model	50
Figure 3.9	Traffic analysis model demonstrating level of automation of the different tasks	53
Figure 3.10	Superimposed sample vehicle trajectories from multiple camera views (2)	54
Figure 3.11	Metadata data model	55
Figure 3.12	Example alignments and links for a typical roundabout	58
Figure 3.13	An example trajectory that jumps from one road user to another	65
Figure 3.14	tvaLib GUI	67
Figure 4.1	Hourly user pairs versus flow rate	75
Figure 4.2	Hypothetical evolution of TTC over time for an ordinary collision course	81
Figure 4.3	Sample size statistics for initial conditions of an example motion pattern array	102
Figure 4.4	Probability map based on the joint arrival of road users for discretized-motion-pattern-based collision course prediction	103
Figure 4.5	Collision point dimension diagram	104
Figure 4.6	Time series of TTC indicators for a single user pair calculated using three different motion prediction methods	105
Figure 4.7	TTC distribution comparison for a before-after or cross-sectional safety continuum comparison	106
Figure 4.8	Serious event comparison of hourly events	107

Figure 4.9	Disaggregated and aggregated time series of all user pairs in test data	108
Figure 4.10	Comparison of hourly serious event counts versus percentage of events	109
Figure 4.11	Cumulative distributions of TTC indicators by aggregation method and prediction method	111
Figure 5.1	Typical four-way roundabout sectioned into component merging zones	118
Figure 5.2	Vehicle-vehicle conflict zones for a typical roundabout merging zone .	120
Figure 5.3	Roundabout merging zone metadata supplied to the HLI roundabout module	122
Figure 5.4	Time instants for calculation of yielding post-encroachment time for a roundabout merging zone	128
Figure 5.5	Roundabout merging zone movement types	129
Figure 5.6	Roundabout locations across the province of Québec	130
Figure 5.7	Basic geometric design elements of a roundabout merging zone	133
Figure 5.8	Sequential arrivals of road users at the approach, and start of all merging zones.	138
Figure 5.9	Sequential arrivals of road users at the approach and start of individual merging zones	139
Figure 5.10	Road user trajectories across the merging zone for two sites of differing angular size	140
Figure 5.11	Speed distributions by lane at a typical multi-lane merging zone . . .	141
Figure 5.12	Spatial distribution of average speed for roundabout merging zones .	142
Figure 5.13	Speed quartiles for each site by movement types (i) through (iv) . . .	143
Figure 5.14	Aggregated mean speed and standard deviation profiles by movement types (i) through (iv)	144
Figure 5.15	Yielding post-encroachment time distributions	145
Figure 5.16	Lead yielding post-encroachment time box plots for each site	146
Figure 5.17	Curve fitting for probability distribution of all TTC values	147
Figure 5.18	Curve fitting for probability distributions of TTC values at top three sites with most user pairs	148
Figure 5.19	Distribution of median angle of incidence per user pair	149
Figure 5.20	15 th centile time-to-collision pooled to six merging zone clusters . . .	151
Figure 5.21	Relationships between traffic characteristics, road design, and traffic safety	165
Figure 6.1	Roundabout locations across Skåne county, Sweden	171
Figure 6.2	Example instalation of a trailer-type video data collection unit	175
Figure 6.3	Sample camera view of a roundabout merging zone	176

Figure 6.4	Distribution of median angle of incidence per user pair	180
Figure 6.5	Comparison of hourly speed averages between Québec and Swedish sites	182
Figure 6.6	Aggregated mean speed and standard deviation profiles by movement types	183
Figure 6.7	Yielding post-encroachment time distributions for Québec and Swedish merging zones	184
Figure 6.8	Comparison of TTC distributions between Quebec and Sweden	188
Figure D.1	Curve fitting of all TTC values for the 41 best probability density function models	226

LIST OF SYMBOLS AND ABBREVIATIONS

- AADT** average annual daily traffic. 37, 135
- CP** collision point. 78, 79, 87–89, 93, 95, 110, 113, 190
- DTW** dynamic time warping. 64, 196
- FCM** fuzzy C-means. 25
- FPS** frames per second. 36, 49, 51
- FRQNT** Fonds de recherche du Québec – Nature et technologies. iv, 5
- FRQS** Fonds de recherche du Québec – Santé. iv, 5
- GIDAS** German In-Depth Accident Study. 2
- GT** gap time. 21, 29, 80, 122
- GUI** graphical user interface. 66, 220
- HLI** high-level interpretation. 56, 57, 121, 124, 136, 181, 220, 221, 223
- HMM** hidden Markov model. 25, 83
- LCSS** longest common subsequence. 25, 64, 85, 99
- MOTA** multiple object tracking accuracy. 14, 51, 59, 60, 65, 189, 193, 219–221
- MOTP** multiple object tracking precision. 51
- MPA** motion pattern array. 59, 70, 85–87, 94
- MTM** Modified Transverse Mercator coordinate system. 47
- MTQ** Ministère des Transports du Québec. iv, 5, 27, 126, 134, 154
- NCHRP** National Cooperative Highway Research Program. 27, 30
- NSERC** Natural Sciences and Engineering Research Council of Canada. iv

- PDO** property damage only. 15, 16, 71, 72
- PET** post-encroachment time. 20, 21, 76, 79, 80, 93, 121
- pPET** predicted post-encroachment time. 21, 70, 79, 80
- RGB** red, green, blue. 44, 45
- RRSR** Réseau de recherche en sécurité routière. iv
- SCC** safety continuum comparison. 81, 82, 96–98, 110, 115, 149, 161, 166, 183, 185, 190, 191
- SEC** serious event comparison. 30, 81, 82, 96–98, 100, 110, 112, 115, 117, 143, 144, 161, 163–165, 185, 190, 192
- SSAM** Surrogate Safety Assessment Model. 22, 29
- SSM** surrogate safety measure. 4, 5, 8, 9, 17, 18, 21, 29, 70–75, 77–80, 82, 83, 85, 92–98, 112–114, 116, 121, 126, 149, 150, 155, 156, 164, 179, 186, 187, 189–191, 193, 196, 198
- TCT** Traffic Conflict Technique. 2, 8, 11, 19, 20, 22, 29, 30, 91, 96, 117, 169, 190
- TRB** Transportation Research Board. 27
- TTC** time-to-collision. 20–22, 25, 70, 79–81, 87, 89, 92, 93, 95–101, 110, 112–114, 116, 121, 125, 136, 144, 146–149, 154, 156, 161–165, 173, 181, 183, 185, 186, 191–194
- UTM** Universal Transverse Mercator coordinate system. 47
- VCU** video data collection unit. 33, 37, 38, 42, 68, 127, 170, 172, 189, 194, 225
- yPET** yielding post-encroachment time. 57, 70, 76, 80, 114, 116, 121–123, 125, 136, 142, 143, 154, 156, 160, 161, 165, 166, 182, 185, 186, 191, 194

GLOSSARY

alignment A spline or series of points in space defining the centre line of a lane, usually associated with a trajectory cluster. 52, 56, 57, 65, 68, 121, 139, 197, 220

analysis zone A polygon defined in world space to which trajectories are constrained for the purpose of performing traffic analysis. 53, 56, 59, 62, 73, 74, 84–86, 94, 100, 112–114, 120, 126, 127, 136, 172, 194, 196, 220

approach A road segment which provides direct access to an intersection, roundabout, or some other road facility. In a general context, approaches typically designate the corridor as a whole, including any lanes enabling exiting of traffic in the direction of flow opposite the approach. In this work, the strictest sense is used, i.e. designating only the portion of the road providing direct access. 26, 29, 73, 76, 117, 118, 120–125, 132, 137, 139, 140, 142, 150, 154, 156, 157, 159–161, 163, 164, 172, 173, 179, 186, 194, 197

braking time Time required by a road user to execute evasive action by braking only. 19, 75

collision A specific type of traffic accident in which two or more moving road users, or a moving road user and a static object, collide with one another as the direct result of transportation activity. 1–5, 8, 15–25, 27–29, 70–75, 77–80, 83, 87–89, 91, 96–98, 110, 113, 114, 116–119, 139, 194, 196, 197

collision course An instantaneous situation where two or more road users have some uncertain, non-zero probability of colliding in the near future. 7, 21, 23–25, 76–80, 82–84, 87, 89, 91–93, 95–99, 113, 114, 121, 144, 161, 162, 190–192, 194, 197

collision point A point in world space where two or more predicted trajectories collide (the road user’s trajectories overlap and their volumes intersect at least once at a common time). 78, 82

collision probability Given a time series of collision courses, this is the probability that at least one collision course of the series results in a collision. The length of the time series establishes the exposure context. 4, 5, 15–20, 22, 25, 29, 71–76, 78, 91, 94, 96–98, 114, 117, 149, 156, 160, 162, 186, 190, 191, 193

- collision severity** Some measure of the total amount of damage done to property and bodily harm done to individuals in the event of a collision, usually measured in a dollar value (property damage or converted from personal injury), or in categories of bodily harm including minor injuries, severe injuries, and death. 4, 15, 16, 18, 19, 22, 27, 28, 71–73, 75, 97, 98, 156, 191
- computer vision** Field of computer science which extracts and analyzes meaningful information from images and video data. 4, 7–9, 11–14, 20, 30, 32, 33, 45, 51, 52, 73, 85, 189, 220
- conflict severity** A term sometimes used in the literature to denote probability of collision based on Traffic Conflict Technique safety indicators, and sometime also collision severity. The term is disfavoured in modernized surrogate safety methods (particularly this one) to avoid terminological confusion with measured collision severity. The more precise terms *collision probability* and *collision severity* are preferred. 16
- conflicting flow** Not to be confused with traffic conflicts, the conflicting flow is the traffic flow corresponding to road users within the roundabout entering the merging zone (as opposed to those users entering the roundabout via the approach). 121–125, 137–140, 142, 160, 163, 197
- constant velocity** A motion prediction method which assumes that no road user control input occurs (e.g. during a collision course). 23–25, 82, 83, 89, 99, 100, 110, 112, 190, 197
- converging manoeuvre** A type of road user interaction where road users' trajectories overlap at a specific point, emerging as a result of the needs and constraints of the geometry, and requiring negotiation to avoid a collision. This situation may or may not involve little to no prior knowledge of each other's presence. 92, 93, 119, 147
- crossing zone** A point in world space where two or more predicted trajectories cross but do not collide. 78, 79, 82, 113, 121
- discretized motion pattern** A proposed novel motion prediction method which learns road user motion pattern probabilities discretely. 85–87, 89, 95, 98–100, 110, 112–114, 116, 144, 161, 163, 183, 191, 194, 196, 197
- dynamic time warping** An algorithm used to determine the similarity of two sequences, e.g. trajectories. 64

- evasive action time** Time taken for a road user to initiate and execute any evasive action (e.g. braking, steering), measured from the instant the action is applied to the end of the collision-course. The passage of time resulting from mechanically distinct evasive actions and processes all serve a similar purpose (e.g. must complete in a period of time shorter than time-to-collision and any dynamically adjusted time-to-collision as the result of ongoing evasive action). 80, 115
- exit** Road segment used to exit an intersection, roundabout, or other road facility. 29, 117, 120, 121, 124, 137, 141, 150, 154, 157, 159, 172, 173, 179, 182, 194
- feature-based tracker** A category of computer vision algorithms which track the motion of features within a video sequence. 12, 14, 42, 43, 45, 50–52, 62, 68, 193
- gap acceptance** Gap acceptance is both a process and a corresponding model used when road users wish to enter or cross a stream of traffic in such a way as to not collide with any approaching road users. Two models exist: a model based on a single gap between leading and lagging road users, used typically for crossing streams of traffic, and a model based on a lead and a lag gap between the merging road user and the leading and lagging road users already present within the destinating stream. 21, 29, 76, 92, 93, 121
- gap time** Gaps are the bumper-to-bumper separations between road users expressed in units of travel time. In surrogate safety, Gap Time has come to mean a particular form of predicted post-encroachment time, defined as “The time lapse between the completion of encroachment by a turning vehicle and the arrival time of a crossing vehicle, if the road users continue with the same speed and path.” (American Association of State Highway and Transportation Officials, 2010). 21, 80
- head-on conflict** The characterisation of a traffic event, whether a user pair, a collision course, a traffic conflict, or a collision, where vehicles originate from opposing directions and approach each other nearly head-on. Note that the term *conflict* in *head-on conflict* should not be confused with the TCT definition of *traffic conflict*. 91–93, 117
- hidden Markov model** A type of statistical model following the Markov process in which states are unobserved discrete random variables and only the model’s output is known and is continuous. It is typically applied to problems which involve recovering sequential data from output of that process, for example: recognizing text (unobserved state) from handwriting (output), or recognizing the intentions of a driver (unobserved state) from the vehicle’s trajectory (output). 25

- high-level interpretation** Analysis of non-generalizable, study-specific traffic parameters. 56, 57, 223
- historical accident analysis** Statistical analysis of historical accident data for the purpose of diagnosing road safety. 1, 4, 8, 19, 28, 71, 72, 119, 139, 174, 187, 196
- historical accident data** Ensemble of accident records generated by either police or ambulance services, or by insurance companies, containing pertinent information about collision history: location, time, type, cause, environmental effects, etc. 1, 2, 15, 16, 18, 19, 70, 72, 94, 98, 114, 116, 119, 168, 174, 179, 181, 187, 191, 193, 196, 198
- homography** A mathematical coordinate transformation between two planes, which projects a point from an image space plane to a world space plane. 46, 47, 56, 62, 220
- image space** The coordinate space used to represent data on the projected plane of the camera's field of view. xxiii, 44–46, 57, 62
- interaction** In general, this term describes the overall relationship between two objects in a user pair, within a distance where trajectory negotiation may be necessary to avoid a collision. In some literature, it can also describe the set of interaction instants and associated measures of safety over the common time interval. 3, 6, 19, 20, 26, 27, 50, 72, 74, 91–94, 96, 114, 116, 119, 120, 123, 126, 139, 147, 148, 162, 163, 173, 191
- interaction instant** An interaction instant is the instantaneous relationship between two objects (e.g. road users) in a user pair, from which collision courses may be defined and surrogate safety measures calculated. 23, 74, 76–79, 82, 87, 88, 91, 93–95, 99, 148, 155, 156, 162, 190
- lane change manoeuvre** A type of road user manoeuvre requiring a displacement into a different lane within a multi-lane traffic corridor (contrast with a converging manoeuvre). As with converging manoeuvres, if other road users are already present in the destinating lane, then a special gap acceptance interaction occurs to allow the road user to enter the lane without colliding with any road users already present. 92, 93, 119
- longest common subsequence** An algorithm used to determine the similarity of two sequences, e.g. trajectories. 25
- mask** A polygon defined in image space for the purpose of constraining tracking efforts or rejecting erroneous trajectories after tracking has occurred. 56, 57, 59, 62, 85, 194, 220

merging zone In the narrowest context, a merging zone refers to a roundabout sub-segment where merging action takes place between approaching and conflicting flows. In a more general context, it refers to the zone proper plus any nearby elements associated with the zone, such as the roundabout approach and exit directly connected to the merging zone. In the broadest context (outside of roundabouts), this is a zone noted for producing, in addition to rear-end conflicts, side-swipe conflicts from converging manoeuvres, and, in the case of multi-lane facilities, from lane change manoeuvres as well. 5, 6, 29, 57, 74, 76, 95, 96, 116–124, 127, 131, 134, 136–141, 144, 147, 148, 150, 155, 157–163, 165–167, 170, 172–174, 179–184, 186, 192, 194, 197, 198

motion prediction The process of generating a set of predicted trajectories from an algorithm emulating the expected (natural) motion of road users at a given instant, based on the road users' positions up to that instant. 7, 23–25, 30, 77, 80, 82–87, 89, 90, 92, 93, 95, 97–100, 110, 112–114, 121, 122, 161, 183, 190, 191, 194, 196, 197

moving object Any object (usually a road user) described by a time series of positions, corresponding velocity vectors, and, optionally, size. 12, 13, 33, 36, 40, 42–45, 50, 51, 56, 60–62, 64, 70, 197, 222

Multiple Object Tracking Accuracy A performance measure (of accuracy) for automated trajectory extraction (e.g. using computer vision). 14, 51

Multiple Object Tracking Precision A performance measure (of precision) for automated trajectory extraction (e.g. using computer vision). 51

normal adaptation A motion prediction method which assumes that only random minimal road user control input occurs (e.g. during a collision course). 24, 83, 89, 99, 100, 110, 112, 190

post-encroachment time A measure of the difference in arrival time at a crossing zone of two road users with overlapping trajectories. 20, 76

predicted post-encroachment time A post-encroachment measure predicted from predicted trajectories instead of observed trajectories, e.g. at a crossing zone. 21, 79

predicted trajectory A trajectory which is generated using motion prediction. 77–79, 82–84, 99

prediction parameter A set of parameters and instructions used to predict a vehicle's trajectory in accordance with a hypothesis of collision precursors and normal driver behaviour. 82, 83

- rear-end conflict** The characterisation of a traffic event, whether a user pair, a collision course, a traffic conflict, or a collision, where vehicles originate from similar directions (e.g. within the same lane). Note that the term *conflict* in *rear end conflict* should not be confused with the TCT definition of *traffic conflict*. 91–93, 117, 119, 147, 148, 162, 186, 191
- road safety** Estimation of collision risk experienced by road users measured in terms of total collisions, collision rate (by some measure of exposure), and collision severity. 2–4, 8, 11, 15, 17, 18, 23, 26–28, 31–34, 69–72, 84, 94–98, 116–119, 123, 131, 138, 149, 150, 158, 164–167, 169, 174, 180, 186, 189, 191, 192, 195
- road safety analysis** Study undertaken for the purpose of estimating road safety and understanding its causes. 1, 2, 4, 6, 18, 19, 23, 33, 37, 68, 96, 112, 113, 115, 116, 148, 189, 191, 193, 196, 198
- road user** An individual operator, human or automated, that makes use of the road infrastructure for transportation purposes and is characterized by independent motion and active decision making. In addition to the operator classification, road users are classified by type of vehicle as well, including: passenger vehicles, trucks, buses, motorized cyclists, cyclists, and pedestrians. This definition, in the context of collision courses, normally does not distinguish between passengers and the vehicle itself. 1, 3–7, 11–33, 36–38, 40, 42, 43, 46, 50, 51, 57, 59, 64, 68–80, 82–87, 89–94, 96–99, 113, 114, 117, 119–126, 131, 136–142, 144, 155, 156, 158–164, 167, 172, 173, 180, 184, 185, 191, 193–197
- road user behaviour** The class of traffic indicators that describe and predict road user behaviour, including: speed, lane changes, gap acceptance, etc. 4–6, 8, 9, 21, 23, 26, 30–32, 56, 76, 77, 80, 82, 84, 89, 116, 117, 121, 149, 155, 160, 166–169, 186, 187, 191–195, 198
- roundabout** A type of circular intersection where through-movements are deflected counter-clockwise around a central island (in right-hand driving jurisdictions). Roundabouts are distinct from traffic circles in that they are designed with orthogonal approaches and exits requiring yielding behaviour on the part of approaching road users. 5, 14, 24, 26–30, 57, 74, 92, 116–121, 123–126, 131, 132, 134–137, 139–141, 144, 154, 155, 157, 162, 164–170, 172–174, 179–184, 186, 187, 189, 191–195, 197, 198
- safety continuum comparison** A surrogate safety method that examines and treats all safety indicators as potential contributors to road safety, and weighs effect by safety

indicator value. 81, 96, 97

safety indicator A measure used to describe the relationship between two road users for the purpose of quantifying collision probability or collision severity in some meaningful way. 5, 7, 8, 20, 22, 23, 25, 70, 77–79, 93–96, 110, 144–146, 149, 161, 165, 183, 190–192

serious event comparison A surrogate safety method that examines only the most serious traffic events, defined by one or several thresholds over one or several safety indicators, sorting events into categories of collision probability or collision severity; less serious traffic events are ignored completely. 3, 30, 96

side-swipe conflict The characterisation of a traffic event, whether a user pair, a collision course, a traffic conflict, or a collision, where vehicles come into close contact by way of a lane change manoeuvre or a converging manoeuvre (e.g. the velocity vectors are not nearly parallel). Note that the term *conflict* in *side swipe conflict* should not be confused with the TCT definition of *traffic conflict*. 92, 93, 117, 119, 147, 148, 162, 186, 191, 198

surrogate safety measure A measure to describe some aspect of road safety, that is not based on historical accident records, e.g. speed. Except in very narrow contexts, this is synonymous with *safety indicator*. 4, 70

surrogate safety method A methodology which applies surrogate safety measures in the analysis of road safety, typically in conjunction with or in lieu of historical accident data. 2, 4–9, 11, 12, 14–16, 18–20, 22–26, 28–30, 32, 70–73, 76–79, 90, 91, 94, 96, 113, 116, 120, 138, 160, 164, 166, 168, 169, 174, 179, 187, 189, 190, 196, 198

time horizon The maximum length of time to predict events (usually trajectories) into the future. 84–88, 98

time series A sequence of measurements made at given time instants and indexed by time. 50, 66, 74, 80, 93, 95, 99, 100

time-to-collision The time remaining for two objects on a collision-course to collide if both road users continue their trajectories on the collision course (i.e. if its outcome is probable and evasive action does not occur). 20, 21, 79, 116

traffic accident A traffic event involving a collision of some kind, or any other undesirable traffic scenario, resulting in property damage or bodily harm. 1, 2, 4, 15, 19

traffic circle A type of circular intersection where through movements are deflected counter-clockwise around a central island (in right-hand driving jurisdictions). Traffic circles are distinct from roundabouts by having tangential approaches and exits, and signalization. at each approach. 26, 154, 165, 192

traffic conflict A traffic event characterized by externally recognizable collision precursor conditions. Two main categories of conceptual definition of a traffic conflict exist in the literature. Many more operational definitions have been put forth and vary greatly from one study to another. As part of this work, this term is relegated to conceptual discussions of problematic road user interactions only; its operational definition is retired in favour of the collision course. 2–4, 11, 19, 20, 22, 23, 27, 29, 71, 73, 76, 96, 116, 119

Traffic Conflict Technique The most common and early surrogate safety method, made of conceptual and operational definitions of traffic conflicts and of a method to estimate road safety from the collected traffic conflict data. 2

trajectory A series of points in space and time defining the movement of an object, such as a road user. Some sources of trajectory data (e.g. computer vision feature extraction) may produce trajectories each composed of the groups of individual features each having their own trajectory. 4, 5, 21, 23–25, 32, 33, 42–44, 50–53, 56, 57, 59, 64–66, 68, 70, 73–79, 82–85, 87, 91, 92, 94, 117, 132, 160, 161, 194, 196, 197

user pair A pair of road users (or one road user and an immobile obstacle) coexisting within a finite distance over a finite common time interval. 73, 74, 76, 78, 80, 87, 91–96, 99–101, 110, 112–115, 120, 126, 144, 146–149, 155, 156, 164, 173, 179, 186, 190, 194, 197, 198

world space The coordinate space used to represent locations of objects in the real world. xxiii, 40, 44, 46–48, 50, 56, 58, 61, 66, 86, 87

yield Behaviour whereby vehicles stop voluntarily (e.g. in the absence of a traffic control device) in front of a merging zone in the event that proceeding through the merging zone cannot be done so safely. This may be imposed explicitly, as in a yield sign, or implicitly, as in textbook rules of priority. In contrast with stop signs, full stops are not required in the event that there is a sufficient gap in the opposing traffic stream to proceed safely. 6, 26, 92, 99, 121, 123, 139, 140, 142, 160, 162, 167

yielding post-encroachment time A post-encroachment time measure taken in the context of expected yielding behaviour. 57, 76, 116, 121

LIST OF APPENDICES

Appendix A	THESIS PUBLICATIONS	216
Appendix B	SOFTWARE	219
Appendix C	STATISTICS	224
Appendix D	COMPLEMENTARY DATA	225

CHAPTER 1 INTRODUCTION

Historical accident analysis, the traditional approach to road safety analysis, makes use of historical accident data, i.e. reported traffic accidents. Traffic accidents, however, are not only undesirable events, they are challenging to collect and limit the potential of road safety analyses, unless public streets are treated as testing grounds for untested road design solutions. Furthermore, traffic accidents can be characterized as chaotic events—in a traffic accident, the difference between life and death may sometimes be measured in milliseconds of reaction time and millimetres of evasion—that have long return periods between occurrences. This makes observation and collection of historical accident data challenging and expensive, as it must be conducted at the government level via transportation agencies and executed via emergency services such as police and ambulance services which compile traffic accident reports.

These reports are not ideal sources of data for the scientific investigation necessary in many traffic engineering projects and research activities¹. The scientific method calls for experimental data to be collected for the purpose of testing a specific hypothesis, not the other way around. That is not to say that most traffic safety studies fit hypothesis to data, but it does at least mean that these sources of data suffer the same issues (i.e. data structure and scope specified and executed by transportation agencies and local authorities) related to any third-party sources of data. They typically lack enough detail or completeness to meet the specific objectives of a scientific study.

Furthermore, these traffic accident reports often present sampling bias, as low-severity traffic accidents tend to be under-reported when authorities are not involved (e.g. fender benders). Finally, comparison between historical accident data in different jurisdictions is challenging due to possible incompatibilities in the level of detail of traffic accidents reporting; incompatibilities in structure or systemic bias (e.g. Bull and Roberts, 1973); and other technical issues such as reporting errors, encoding discrepancies, and reporting incompleteness. Reliance on historical accident data poses challenges in the field of road safety analysis and traffic engineering in several ways, namely:

¹Terminological note: Traffic accident reports completed by police and ambulance services may include scenarios falling outside the purview of traffic engineering (e.g. delinquency). For the purposes of this work, the term *collisions* is used uniquely to identify the specific group of traffic accidents involving the failure of one or more road user to use a road facility appropriately and under normal operating conditions, as these are treatable via traffic engineering methods. Reported traffic accidents which do not follow the definition of collision fall outside the scope of this work.

- This reactive design process means that road safety is typically addressed many years after traffic accidents have occurred and been recorded. New designs must be tested in live scenarios with unknown or only hypothetical dangers. Alternatively, new designs are implemented slowly in the form of pilot projects to minimize exposure to unknown danger with, however, significant delays, as accident data is collected slowly, over decades, or otherwise in very limited quantities.
- As stated previously, accident reporting must be conducted at a regional level, sometimes with the involvement of multiple different agencies or transportation authorities. Arguably, this data needs to be collected in any case for the purpose of national reporting and insurance purposes, however the reuse of pre-collected data to test new hypotheses is not scientifically rigorous and may cause sampling issues. Regenerating the equivalent historical accident data at the level of detail and scope necessary to answer a hypothesis is prohibitively expensive. Some specialised efforts to address this issue with historical accident data have been attempted, such as in the German In-Depth Accident Study (GIDAS), but are otherwise almost never observed in practice.
- Reporting bias can hinder study objectives. Quality of traffic accident reporting depends largely on the degree of involvement, if any, of different authorities (e.g. police versus ambulance intervention, Bull and Roberts, 1973). Alternatively, different authorities may implement different levels of reporting completeness. In one example, a recent road safety analysis studying a lane-change ban on urban highway ramps was to be initially completed with police accident reports, however these had no geographical information other than a highway number and direction, and not always even that. The study had to pursue alternative methods of road safety analysis, including ambulance records (St-Aubin, 2012).

In the 1960's, researchers at General Motors (Perkins and Harris, 1968) attempted to tackle these issues by defining procedures for predicting collisions from ordinary samples of traffic interactions instead of waiting for and observing traffic accident events directly. The general procedure, known at the time as simply *conflict analysis*—now today classified as Traffic Conflict Technique (TCT) and more broadly categorized as a subset of surrogate safety methods—garnered interest quickly. Although no reference implementation exists, the general principle behind TCTs is the observation and measurement of traffic conflict.

In these frameworks, there exists two major competing definitions of traffic conflicts (Zheng et al., 2014a). The first one defines traffic conflicts as discrete, observable events “involving

two or more road users², in which the action of one road user causes the other road user to make an evasive maneuver to avoid a collision” (Parker and Zegeer, 1989). The second school of thought, building on the idea that traffic events or traffic interactions lie on a continuum defined by some measure of proximity, proposes that the most serious events be classified as traffic conflicts according to some proposed threshold (Amundsen and Hydén, 1977).

Adoption of the methodology soon spread to England (Spicer, 1973; Grayson, 1984), Israel (Hakkert, 1984), The Netherlands (Kraay et al., 2013, , original publication in 1986), Sweden (Hydén, 1987), the United States (Parker and Zegeer, 1989), Canada (Brown, 1994), and elsewhere. However, the methodology was criticized on the basis of several unresolved issues (Williams, 1981; Krusysse, 1991; Chin and Quek, 1997), namely:

- Subjectivity of observations: Most of the methodology developed involved counting traffic conflicts according to purpose-defined conditions of safety, ranging from “gut feeling”, to systematic visual estimations of speed and times to accident (Almqvist and Ekman, 2001) by manual observation of traffic by individual observers. Several instructional manuals were produced (e.g. Amundsen and Hydén, 1977; Parker and Zegeer, 1989; Almqvist and Ekman, 2001) for the purpose of training individual observers. Still, criticism of subjectivity in observations remained, especially concerning the limited ability of humans to estimate speed and position of road users accurately in real-time³, the emphasis on using “gut feeling” of safety with some methods, and differences in ability and judgement between individual observers (Chin and Quek, 1997).
- Traffic conflict definition: Not only did different schools of thought exist regarding traffic conflicts, different studies would use different criteria for traffic conflict (Williams, 1981). Various measures were proposed in an attempt to quantify the impact of traffic conflict events on road safety or produce a definition that distinguished dangerous traffic situations from benign situations (the *serious event comparison* approach). However, this resulted in many incompatible measures and definitions being generated, often with simplified thresholds of admissibility that would discard a large number of events

²Note that in this context, and for the remainder of this work, road user denotes an individual operator, human or automated, that makes use of the road infrastructure for transportation purposes and is characterized by independent motion and active decision making. In addition to the operator classification, road users may be classified by type of vehicle as well, including passenger vehicles, trucks, buses, motorized cyclists, cyclists, and pedestrians. This definition normally does not distinguish between passengers and the vehicle or operator.

³Speed estimations performed by humans are surprisingly accurate (as demonstrated in Svensson, 1998), but still less than purpose-built instruments.

without considering the impact or collision probability of individual events, though others have had some degree of success calibrating traffic conflict (e.g. Hydén, 1987).

- Transferability of results: It is still unclear how results from traffic conflict studies defined in one study can be compared with results from traffic conflict defined using a different approach in a different study, especially if the exposure to traffic conflicts is not constant. This is particularly important when design not only affects road user behaviour, but also affects the method of measuring the traffic conflict itself, as several competing and often incompatible traffic conflict definitions exist, or when observations are biased with subjectivity (Williams, 1981).

Furthermore, while the relationship between observed traffic conflicts and the expected number of traffic accidents, expressed as “conversion factors”, has been validated to some degree (e.g. Hydén, 1987), this has only been conducted successfully for a handful of traffic conflict methods, and only for specific studies. It remains to be seen if the same predictive power can be found in other studies. At this stage, the only consensus in the literature is that of a qualitative relationship between traffic conflict and expected number of traffic accidents (or other aspects of road safety).

Issues with subjectivity in observations have by now largely been addressed, thanks to the automation of road user motion measurement and a more complete framework for defining and measuring traffic conflicts, or even surpassing the need for such a framework conceptually. This has been made possible in part thanks to advances in computer vision which enable the automated, and thus consistent and repeatable, observation of road user trajectories from video data. With detailed recorded road user trajectories and a robust road safety framework, any type of surrogate safety measure (SSM) can then be prototyped, evaluated, systematically compared (since the same test data can be used across different methodologies), and eventually validated. With this framework, some empirical validation and statistical regression, and sufficient data needed to address model validation, surrogate safety methods promise to upend, or at the very least complement, historical accident analysis, assuming SSMs with significant and generalizable collision prediction power exist. Still, a paradigm shift in road safety analysis is needed.

In the meantime, however, a few specific SSMs are already being used, the prime example being speed. Speed has a significant body of research behind it that has consistently demonstrated the link between it and collision severity. The relationship between operating speed and collision probability has also been studied extensively, though the link between these

two is more complicated and some debate regarding this relationship persists (Hauer, 2009). This can be illustrated with a simple example: on average, while highways produce more severe road-user-to-road-user collisions than do local roads (e.g. Elvik, 2009; Institute for Road Safety Research, 2012), especially if controlling for the same types of road users, the observed rate of collisions, i.e. collision probability, between road users does not necessarily follow this trend.

The issue is further compounded by the fact that multiple definitions of collision probability exist (contrast collisions per million vehicles, collisions per vehicle-kilometre travelled, and collisions per million trips). Furthermore, exposure with, for example, traffic has been demonstrated as being a factor which influences individual collision probability—that is to say, the relationship between the total expected number of collisions is not linearly proportional to exposure (Qin et al., 2004; PIARC Technical Committee on Road Safety, 2004). While exposure is observed as an important contributing factor capable of predicting total expected number of collisions, the “safety in numbers” effect generally dictates that, in many cases, collision rates (collisions per individual, i.e. collision probability) decrease as traffic exposure increases (Brüde and Larsson, 1993; Leden et al., 2000; Jacobsen, 2003).

This research project aims to (i) tackle some of the aforementioned problems of feasibility of surrogate safety methods by developing the tools necessary for the systematic, large-scale extraction of road user trajectory data; (ii) outline a generalized framework for SSMS and other safety indicator identification and comparison approaches; and (iii) improve surrogate safety methods and conduct an exploratory analysis of these compared to older strategies. This research project then examines the application of the methodology in the context of a road safety research project of 35 roundabout merging zones in Québec, followed by a comparison of road user behaviour between 10 roundabout merging zones in Sweden and an equivalent sample of merging zones in Québec.

Although the video data collection system and surrogate safety methods are designed and tested to be applicable to any type of traffic scenario⁴, there are several reasons why a study of roundabouts was chosen to illustrate application of the methodology, including:

- There is a limited number of road safety studies specific to roundabouts performed in Canada, and, to a lesser extent, the United States.
- This prompted the Québec Roundabout Safety research project (proposal number 2012-SO-163493 funded by the Fonds de recherche du Québec – Nature et technologies (FRQNT), the Fonds de recherche du Québec – Santé (FRQS), and the Ministère des

⁴Case-specific traffic analysis can be handled as well—a reference example is presented in Chapter 5.

Transports du Québec (MTQ) (Saunier et al., 2015; St-Aubin et al., 2013a; Burns et al., 2013)) making funding and data readily available for the development of an improved road safety analysis methodology.

- Roundabouts are infrastructures without explicit stop control; as such all interactions are managed exclusively by road users themselves through yielding action, thus making the study of roundabout road user behaviour particularly relevant for studying interactions.
- The complexity of naturalistic vehicular motion through roundabout merging zones lends itself particularly well to the development and testing of complex surrogate safety models.
- For international comparison—a feat never undertaken at this level of detail—roundabouts constitute one, if not the best, type of road infrastructure for controlled direct comparison, as road design is nearly identical for European and North American roundabouts, unlike other forms of road infrastructure. This is because the design in North America has been imported directly from Europe (Rodegerdts et al., 2010).

1.1 Research Objectives

The objectives of this research project are multifaceted and interlinked. This project aims to:

- build a framework for a complete, automated, large-scale, microscopically-detailed, video-based traffic data collection and analysis system in order to
- improve proactive road safety analysis via objectively and systematically defined surrogate safety methods, such that a comprehensive, detailed, and large-in-scope road safety analysis can be made,
- provide an illustration of the above by means of a case study of North American roundabout design using a sample of Québec roundabouts, and
- quantify design from an international perspective, using an additional sample of Swedish roundabouts.

Each of these four objectives is divided into its own chapter and ordered in such a way that the reader may follow the completion of the research objectives in a linear fashion (e.g. completion

of a systematic surrogate safety method first requires an appropriate data collection tool), though some amount of iteration is present.

1.1.1 Contributions

More specifically, the contributions of this research are as follows (in order of appearance):

- Improvement of existing computer-vision-based automated traffic data collection, and analysis tools and procedures, including:
 - the development of equipment, techniques, and procedure for large-scale video data collection efforts;
 - practical solutions to the challenges of automated processing of large quantities of video and trajectory data (big data), and issues with unsupervised or partially supervised analysis;
 - issues related to the quality of computer-vision-based trajectory data extracted from video data: video data processing and optical error correction, tracking performance optimization, and tracker error correction using filtering functions.

The contribution in this area is thus composed of many small, practical, and theoretical advances necessary before a large-scale, microscopic, computer-vision-based traffic data collection system can be assembled, and vital to the remaining objectives. To this end, a shared-source⁵ software library, *tvaLib*, is developed to implement the contributions contained in this work.

- Development of a systematic surrogate safety method into a workable state:
 - an expanded surrogate safety method structure to include classical and new collision course prediction methods, new safety indicator aggregation methods, measures of exposure, and methods of evaluation;
 - an improved road user motion prediction method for non-linear motion environments (as is the case in, for example, roundabouts); and
 - an exploratory analysis of surrogate safety methods performance, and applicability (no conversion factors at this stage, however).

⁵At the time of writing, *tvaLib* does not meet the Open Source Initiative's definition of *open source* and is thus designated as *shared-source* software by the OSI instead. Regardless, the source code is freely available and usable for purposes of research.

- A case study of road user behaviour, road safety, and traffic flow of Québec roundabouts intended to explain aspects of road safety from specific elements of road geometry design and land use parameters. This serves both as a contribution in and of itself and as a demonstrative application of the large-scale, microscopic, computer-vision-based traffic data collection system and new surrogate safety method framework. This case study is conducted with a sample of approximately 3.7 % of the entire population of registered drivers of Québec.
- An international comparison of roundabout safety measured at a microscopic-level of detail between similarly designed and similarly situated (relative to the built environment) roundabouts in Québec and Sweden. A comparison at this scale has never been undertaken before and should offer further insight into the differences in driving culture between North American and European drivers. This is especially possible, as roundabouts offer some of the most comparable road infrastructure designs between the two continents.

Some of these contributions are visible in the document outline of Figure 1.1.

1.1.2 Limitations and Scope

This project does not cover specific implementations of computer vision tracking algorithms. While continually evolving, computer vision is a mature field within computer science. This project is built upon OpenCV-based feature tracking algorithms (Brahmbhatt, 2013) as implemented for traffic analysis in the *Traffic Intelligence* project (Saunier and Sayed, 2006a; Jackson et al., 2013).

This project is not intended to be a detailed validation of surrogate safety methods as a whole or of any specific SSMS. For one, validation in the form of conversion factors has already been completed to some extent using TCT safety indicators (Hydén, 1987). These are arguably less reliable than existing SSMS such as speed (Pasanen, 1992; Rosén and Sander, 2009) and are, arguably, less reliable than the more systematically-defined SSMS (to be reviewed and revisited as part of this work). Secondly, and more importantly, the project scope necessary to complete a thorough validation of SSMS is beyond the scope of this, or any similarly-sized, project. The creation of further conversion factors and methodological calibration for specific SSMS requires considerable resources, including a purpose-built dual data collection program collecting both traffic collisions data and microscopic traffic data (i.e. video data). Fortunately, this is in the research mandate of an ongoing research project (Laureshyn et al., 2015). Stated differently, the very challenges of historical accident analysis that surrogate

safety methods attempt to address are slowing the research progress and development of surrogate safety methods. However, this issue can be addressed with appropriate funding, and the data need only be collected a small number of times for the method to be properly calibrated.

1.2 Outline

This document is organized into thematic chapters, as demonstrated in Figure 1.1. A state-of-the-art of the scientific and technical literature regarding computer vision applied to traffic analysis, SSMs, and roundabouts is presented in Chapter 2. Chapter 3 presents the approach and methods for traffic analysis based on video data, including data collection, processing, and trajectory extraction. This chapter also contains a small section presenting tracking accuracy calibration results necessary for the remaining chapters. Chapter 4 documents the SSM framework, and includes some data analysis as part of the guiding choices in methodology selection. Overall, these two chapters are methodology- and theory-oriented.

Chapters 5 and 6 present the case study of roundabouts, with respect to Québec (North American) roundabout design and road user behaviour, and an international comparison between Québec and Swedish roundabout road user behaviour, respectively. Chapter 7 closes the discussion with concluding remarks, contributions, limitations, and future work.

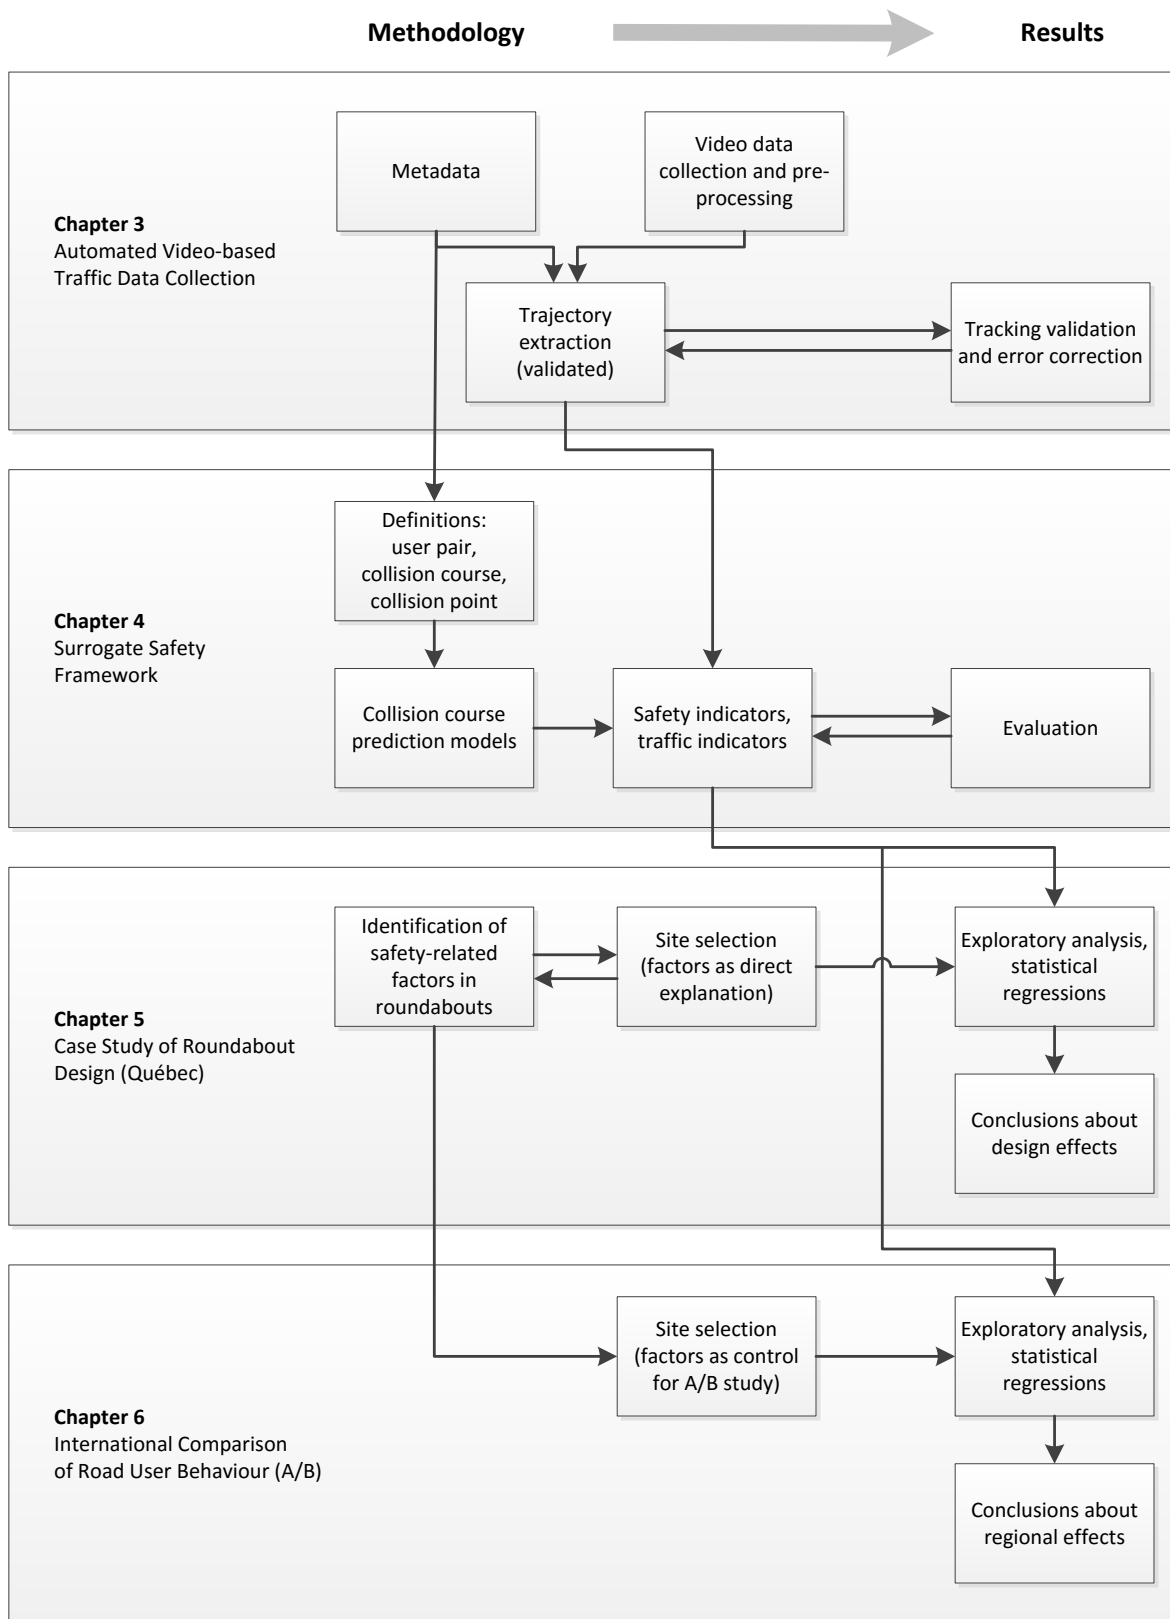


Figure 1.1 Outline of the overall proposed methodology and results. Each major topic is contained in one chapter.

CHAPTER 2 LITERATURE REVIEW

This chapter is intended to provide the uninitiated reader with a brief historical background, a primer on the state of the art of computer vision applied to traffic data collection and analysis, and a basic road safety framework from which to re-explore developments in the field of surrogate safety methods. It concludes with a discussion of roundabout design, roundabout road safety, and any relevant roundabout studies implementing surrogate safety methods.

2.1 Video-Based Traffic Data Collection Systems

Data collection is a fundamental challenge and activity in any type of traffic study, and is arguably the most arduous process of all, given the long periods of observation required and increasingly large networks or locations managed. A variety of approaches have been formulated for a wide range of traffic data collection needs, ranging from basic aggregated road user counts and speeds, to more sophisticated measures such as traffic conflicts. The data collection systems used in the past, and still today, range from the rudimentary pen and paper to various embedded sensors, each with their advantages and disadvantages.

Some of the more traditional classes of embedded sensors include magnetic induction (inductive loop detectors), radar, and visible light sensors (cameras), to name a few. Today, traffic cameras are ubiquitous, despite minor challenges presented by lighting conditions, and the use of computer vision is gaining traction as a feasible method of automatically collecting traffic data using traffic cameras as sensors. Not only is video data easily obtainable, it is one of the only methods capable of producing effectively *continuous* spatial measurements (30 observations per second or more) of all distinct moving objects such as road users within a scene passively. Although Global Positioning Systems do provide trajectories, these systems require the active participation of the road users being monitored and are generally less accurate (Lou et al., 2009). But use of video data as a form of traffic data collection predates modern applications of computer vision. In this approach, video data is simply used to more conveniently conduct manual traffic counts or, in the case of TCTs, traffic conflict observations. Meanwhile, manual analysis is still conducted on video data for even the most sophisticated systems, as a form of ground truth data collection to validate automated video data collection systems (Jodoin et al., 2016) or manual surrogate safety methods (Laureshyn, 2013).

As a traffic data collection tool, video-based automated traffic data collection has been made

possible with the advent of, and advancements in, the field of computer vision, and through decreasing costs in video sensor technology. The most basic computer vision systems act as simple area detectors, detecting the presence of road users in a pre-defined region of a given scene. The sophistication of implementation ranges from this basic binary presence detection (“on/off” presence of a “virtual loop”¹) (Michalopoulos, 1991) to high-resolution continuous motion tracking of distinct objects in real-time (Coifman et al., 1998).

The former variety of sensors are better geared to practitioners looking for simple traffic volume counts, given the simplicity and reliability of these systems, while the latter are better suited for researchers and practitioners looking to undertake more complex data collection in behavioural, safety, and road infraction studies, among others. For the purposes of the surrogate safety methods discussed later, the remainder of this work will examine and implement computer vision of the high-resolution variety. However, real-time processing is not a requirement (and is, in fact, discouraged if such a system detracts from system affordability and tracking accuracy), so long as data processing times are reasonable and keep pace with data collection speeds.

2.1.1 Overview of Automated Computer Vision Systems

Buch *et al.* published a review of computer vision techniques, and their limitations, for vehicle tracking (Buch et al., 2011). Two major families of motion tracking exist in computer vision for stationary cameras:

- Feature-based tracking, which tracks individual moving features from a scene using a specific feature-based-tracking algorithm, the **feature-based tracker**. Conceptually, features are very small groups of pixels with sufficient contrast and regularity to continuously identify the smooth motion of a point on a moving object from one video frame to the next. Then, those features with similar motion, and thus representing the same physical object in motion, are clustered together to form an entity known as a moving object. The main limitation of feature-based tracking is that it is sensitive to the appearance and disappearance of individual features present on the object being tracked. This usually occurs during partial occlusion (even very slight or momentary) of the object, rotation of the object in three dimensions such that the surface becomes hidden, or when the resolution of the object changes and details emerge or vanish, i.e.

¹The “virtual loop” refers to the virtualization of a loop detector, which in turn is a prevalent form of traditional traffic data collection which measures the passage (“on/off” states) of a vehicle directly above the detector via magnetic induction. This technology has been deployed widely in commercial applications (e.g. Autoscope).

when travelling to or from the horizon. Furthermore, feature tracking must rely on motion similarity exclusively to group features into moving objects.

- **Background subtraction**, which creates a model of the background (stationary elements in the scene), is an alternative tracking algorithm. The most common implementation calculates an averaged background image from which to ignore pixels of similar colour on subsequent frames. Remaining pixels normally form blobs representing moving objects within the scene of a colour normally different from that of the background. Background subtraction can take advantage of the significant contrast between pavement, which is typically some monochrome shade of grey, and vehicles, which are mostly painted in various bright colours or with non-monochrome features, e.g. red tail-lights. However, the algorithm is subject to limitations with contrast between the background and the object, and is particularly sensitive to lighting conditions. Some example applications of background subtraction used in traffic applications include a crowd detection algorithm (Maurin et al., 2005), and (Milla et al., 2010) which published an entire chapter on the subject, among others.

Some additional tracking algorithm-independent challenges can be highlighted, particularly regarding object occlusions from large vehicles such as trucks or buses, or other, tall, road-side obstacles, and difficulty dissociating subjects travelling at similar speeds in close proximity to one another. Finally, parallax is a concern for especially low-angled video data. Fortunately, most of these challenges are expected to be overcome as volumetric recognition techniques or more advanced artificial intelligence scene recognition techniques emerge such as particle filtering-based tracking-by-detection (Breitenstein et al., 2011). Alternatively, these challenges can be overcome with better overhead video data collection, such as is the case with increasingly popular drones (Puri, 2005).

Some of the earliest practical forays into applied computer vision date back to 1998 (Coifman et al., 1998). Coifman *et al.* listed five criteria of data collection effectiveness: (i) automatic segmentation of each vehicle; (ii) correct detection of all types of road vehicles; (iii) function under a wide range of traffic conditions; (iv) function under a wide variety of lighting conditions; and (v) operation in real-time. Criteria (i), (iii), and (iv) are obviously necessary. However, for research purposes, real-time operation is not necessary at all, and can be detrimental if the equipment required reduces accuracy or reliability, or increases cost. Arguably, road user classification is not strictly necessary either, as will be seen in the discussion of the methodology and the case studies for this project, though it is valuable data in and of itself and could be helpful in boosting the performance of some of the analyses.

Additional notable applications of computer vision to video-based traffic data collection tools include the NGSIM project, which produced a dataset of extracted trajectories from video data of four corridors composed of freeways and urban arterials each covering several kilometres using overlapping cameras (Kim et al., 2005); the SAVEME project which fielded a small but early implementation of video tracking for use in a surrogate safety method framework (Ervin et al., 2000); and the SHRP 2 Site Observer system (Gordon et al., 2012) which presented a complete automated video-based traffic data collection system, though not as capable of being implemented on a scale as large as proposed in this project.

Estimation of data quality was performed on NGSIM data and revealed particular issues with acceleration measurements that were susceptible to noise (Punzo et al., 2009). A general criterion to estimate unbiased trajectories was put forth. Since then, a more thorough measure of tracking accuracy—the Multiple Object Tracking Accuracy (MOTA)—has been developed and will be used for this project (Bernardin and Stiefelhagen, 2008). MOTA is a measure of accuracy that is noteworthy for evaluating *all* tracks *continuously*, instead of point-based accuracy estimates common in the literature.

Today, *Traffic Intelligence* (Saunier and Sayed, 2006a; Jackson et al., 2013) is probably the leading initiative in the application of computer vision to video-based traffic data collection. It is based on the massively popular, open-source, and robust OpenCV (Brahmbhatt, 2013) computer vision framework, implementing the Kanade-Lucas-Tomasi feature-based tracker (Tomasi and Kanade, 1991). The primary contributions of this project are a readily available feature-based tracker for traffic data collection, a trajectory data storage scheme, and a basic and generic surrogate safety method framework. However, this project is still not practice-ready, as it is not designed for processing large volumes of data, is relatively complex to operate, and lacks task automation and self-regulation features.

Meanwhile, road user classification is an active area of research in the domain of computer vision applied to traffic analysis. It includes many different algorithms, notably object classifiers (Breitenstein et al., 2011). This task is beyond the scope of this work, as traffic movements found in the case-analysis (roundabouts) can be easily assumed to be heterogeneous (motorists, cyclists, and pedestrians do not share lanes and travel orientations as is the case in other types of infrastructure). However, for future projects, recent advancements in classification can be easily integrated into the work flow of analysis to integrate classification on the basis of frequency analysis of speed profiles (Saunier et al., 2011a) and appearance (Zangenehpour et al., 2015) into Traffic Intelligence.

2.2 Safety Framework

Before surrogate safety methods can be addressed, a review of the fundamentals of road safety theory is necessary. Many studies use the term road safety liberally, especially in statements of study contributions. The problem is that the term *road safety* encompasses many independent and distinct concepts affecting road safety. For example, some types of traffic collisions are severe but infrequent, while others are not serious but may occur frequently. Thus several competing metrics of road safety can exist: expected number of accidents (estimated from historical accident data), property damage, rate of injuries, or fatalities. This list might arguably also include issues related to perceived safety, though these are more difficult to quantify. In any case, suggesting that road safety is improved after the implementation of some road treatment implies that all aspects of road safety have improved—or at least that none of the aspects have deteriorated—which is not always the case. An increased *sense* of danger may be beneficial to reduce collisions, or it may be a criterion of its own. This confusion can be compounded by the ambiguity between the terms *traffic accident* and *collision*. In this chapter and subsequent chapters, the term collision is used exclusively and unambiguously to denote any type of recorded traffic accident involving one or more road users that is directly related to transportation activities and thus treatable via traffic engineering methods or some other direct method of affecting road user behaviour of a transportation nature.

Ultimately, road safety is a risk management activity. Classic risk theory proposes that risk is the “effect of uncertainty on objectives” (ISO 31000, 2009). As such, a collision risk management model for road safety (van Poortvliet, 1999; Hakkert et al., 2002; Saunier, 2005) can be derived using the following analogous components:

- The *effect* of the collisions itself is measured as the **collision severity** which predicts or observes the impact and damage of a collision once it unfolds. This can be measured in one of three ways: property damage only (PDO), injury, or fatality. Injury and fatality classifications have been standardized via the Injury Severity Score (Baker et al., 1974). Some practitioners convert injuries and fatalities into equivalent PDO (as in National Safety Council, 2013, for example), though these are non-standard applications.
- The *uncertainty* is that of road users colliding—measured as the **collision probabil-**

ity²—predicting, or observing the probability, that a collision of any kind occurs for a given traffic event, or alternatively, an expected number of collisions given a number of traffic events, i.e. collision **exposure**. Note that how those events are defined is of great importance.

- The *objective* is a reduction of either collision severity, or collision probability, or both.

Given the existence of a variety of collision severity measures and the different ways of defining traffic exposure, a number of objective measures of collision risk can be developed. In the most classic model, using PDO as an example metric of collision severity, *expected PDO p.a.* (the risk to be minimized) can be modelled from a known collision probability, i.e. a rate of collision measured from historical accident data collected over a given time period (e.g. *p.a.*) and average collision severity of those collision as measured in units of PDO, usually normalized cost (e.g. National Safety Council, 2013):

$$\textit{expected PDO p.a.} = \textit{collisions p.a.} \times \textit{average PDO} \quad (2.1)$$

However, since historical accident data, and thus measures of collision rate, are not available with surrogate safety methods (by design), an alternative form of collision probability needs to be designed. Thus, in modelling collision probability with a surrogate safety method, the scope of exposure is redefined.

2.2.1 Collision Probability

While collision severity is typically measured in terms of PDO, injury, and fatality, collision probability is usually measured as an expected number of collisions for a given amount of exposure. The choice of exposure depends on the context. In a general context, annual exposure is used (to normalize data macroscopically), such as is the case with the measure of *expected number of collisions p.a.*

Archer provides a popular and robust definition of exposure: “Measure of spatial or temporal duration in the traffic system in relation to the number of dynamic system objects road users, vehicles (axles), etc.” (Archer, 2004)

Specific exposure measures are used popularly throughout the literature, resulting in various

²Terminological note: historically, the term *conflict severity* has been used in certain surrogate safety methods to describe collision probability, but this use is disfavoured in modern surrogate safety methods as the term is ambiguous with the term collision severity.

normalised road safety statistics such as: the expected number of collisions per million-vehicles, the expected number of collisions per passenger, the expected number of collisions per vehicle-km travelled, and so forth, or possibly even a combination of two or more of these, e.g. passenger-km. The choice of which measure to use may depend on the transportation objectives—does the specific issue concern moving large volumes of people or travelling greater distances?

All else being equal, as exposure increases, the expected number of collisions also increases. However, the expected number of collisions per unit of traffic exposure does not necessarily increase proportionally, as has been demonstrated with the “safety in numbers” effect. This hypothesis dictates that while a positive relationship exists between the expected number of collisions per unit of time and traffic exposure, the relationship is generally not linear for a number of different traffic scenarios (Brüde and Larsson, 1993; Leden et al., 2000; Jacobsen, 2003). Stated more simply, as annual traffic volume of a group of road users increases, so does the annual expected number of collisions, however the collision probability of each individual member of that group does not increase proportionally, in fact it may, and often does, decrease.

The common explanation of the “safety in numbers” effect is that road users dynamically adjust their behaviour to account for exposure, that the movement of larger groups is more readily visible and predictable than smaller ones or individuals, and that “collective vigilance” (Bhatia and Wier, 2011) may be performed by pedestrians. This has been observed more generally in studies of herd mentality. A classic example of safety in numbers is right-turning-motorists in conflict with crossing pedestrians: the motorists are more likely to be aware of the presence of pedestrians and each additional pedestrian present increases the safety buffer for a substantial number of other pedestrians. Note however that this effect only means that the collision probability of each individual road user tends to decrease, not that the total number of collisions decreases over time. In fact, in most scenarios where “safety in numbers” has been observed, the total number of collisions still tends to increase, only at a decreasing rate, when other measures of exposure, such as traffic volume, increase. The implication is that exposure is effectively a type of very crude SSM capable of modelling the *expected number of collisions*, though non-linearly. Some researchers (e.g. Bhatia and Wier, 2011) remain critical of some of the “safety in numbers” causation hypotheses first proposed (Jacobsen, 2003). Nevertheless, the evidence clearly demonstrates the relationship. In all likelihood, a complex variety of specific causal factors is at play, including road user environmental safety awareness due to a more diversified modal share, and better advocacy for vulnerable users (Bhatia and Wier, 2011).

2.2.2 Collision Severity

Besides modelling the *expected number of collisions p.a.*, the alternative use of surrogate safety methods is modelling collision severity.

In a traffic collision, the primary cause of damage can be attributed to the mostly inelastic release by the objects involved of kinetic energy maintained immediately prior to impact (Fildes and Lee, 1993; Elvik et al., 2004). A fraction of this energy is conserved as kinetic energy (if) the colliding road users, or a single road user and a previously fixed object, continue to move after the instant of impact before coming to complete rest, but a large portion of it is usually absorbed as inelastic losses, via deformation and mechanical failure. This energy is dissipated very quickly, causing impact forces powerful enough to damage vehicles, property, and occupants (road users). How this inelastic energy is dissipated depends on a number of factors: individual characteristics of road users, i.e. susceptibility to injury; vehicle safety performance; road infrastructure features designed to absorb collisions; and some element of chance. But the principle of conservation of mass dictates that some important amount of inelastic energy will be dissipated as deformation (damage) of the colliding bodies. The total amount of kinetic energy available for release can be simply calculated using the absolute velocity with respect to the resting road surface (v_{abs_i}) and vehicle mass m_i for vehicle i :

$$E_k = \sum_i \frac{m_i v_{abs_i}^2}{2} \quad (2.2)$$

It thus follows that v_{abs} and m_i are important factors in road safety analysis. And in fact, driving speed is a particularly relevant issue in innumerable road safety research papers and a hot topic issue with respect to speed limits. Use of driving speed in lieu of historical accident data, sometimes implicitly, in research papers effectively places driving speed as an important and well understood and researched SSM for predicting collision severity in the literature.

It might be argued that the “safety in numbers” effect (Brüde and Larsson, 1993; Leden et al., 2000; Jacobsen, 2003) not only plays a role in collision probability, but may also play a role in collision severity simply due to the nature of congestion. Consider the scenario in which, as the number of vehicles on a road increases, the operating speed tends to decrease along with collision severity due to congestion (Massie and Campbell, 1993).

2.3 Surrogate Safety Methods

Surrogate safety methods are a means to facilitate proactive road safety analysis from field observations, by minimizing data collection costs and addressing issues with small historical accident data sample sizes and durations, and public exposure of untested road designs. It has also been argued that surrogate safety methods provide more context-appropriate information, particularly with regards to providing an understanding of the relevant underlying collision and failure mechanisms (Tarko et al., 2009). The general principle behind surrogate safety methods is that road user interactions follow a hierarchy ranging from isolated passage (no interaction or proximity whatsoever) to collision. This is usually illustrated with a pyramid in the literature (attributed to Hydén, 1987).

2.3.1 History and Purpose

Surrogate safety methods are intended to address a fundamental problem with the traditional historical accident analysis approach for conducting road safety analysis which is that traffic accidents are real, observed events with long return periods and important consequences. Surrogate safety methods make use of the indirect observation of collisions via collision precursors, by way of observation of ordinary driving conditions—much smaller and more efficient to collect than historical accident data.

The mainstay of surrogate safety methods is arguably motor vehicle speed, following the mechanisms outlined in section 2.2.2. It is the subject of countless studies relating collision severity to speed (Pasanen, 1992; Rosén and Sander, 2009); at least one full framework (Elvik et al., 2004); and, to a somewhat lesser extent, collision probability predicted from speed (NHTSA, 1997; Shinar, 1998; Hauer, 2009). For example, it can be clearly demonstrated that speed affects operational characteristics such as braking time, and can account for the increased number of collisions observed in adverse weather conditions, where braking time is significantly longer for higher speeds (Pisano et al., 2008). In turn, many more studies have examined the relationship between speed and every human, environmental, road design, and driving factor imaginable.

TCT was the first development of surrogate safety methods. The earliest form, simply called “conflict analysis”, was originally conceived at General Motors in the 60’s (Perkins and Harris, 1968). The principle behind the TCT is the observation and evaluation of traffic conflicts. In this framework, traffic conflicts were originally defined as discrete, observable events “involving two or more road users, in which the action of one road user causes the other road user to make an evasive maneuver to avoid a collision” (Parker and Zegeer, 1989). Methods

similar to this TCT were developed throughout the world, i.e. in England (Spicer, 1973; Grayson, 1984), Israel (Hakkert, 1984), Sweden (Hydén, 1987), the United States (Parker and Zegeer, 1989), Canada (Brown, 1994), and elsewhere in the 70's and 80's.

An additional school of thought emerged, building on the idea that traffic events or traffic interactions lie on a continuum defined by some measure of proximity between road users. It proposed that the most serious events be classified as traffic conflicts according to some pre-defined thresholds (Amundsen and Hydén, 1977).

However, early surrogate safety methods presented problems of their own, including cost, subjectivity, problematic reproducibility, and non-transferability (Williams, 1981; Krusysse, 1991; Chin and Quek, 1997). Counterarguments to issues of subjectivity have been put forward on the finding of a surprisingly reliable capacity of traffic observers to estimate speed (e.g. Svensson, 1998) and the consistency of subjective surrogate safety method results between observers (Shinar, 1984), although it should be pointed out that the benchmark objective traffic conflict measure used at the time was very rudimentary, i.e. a “weighted deceleration”. Validation of traffic conflicts as substitutes of *expected number of collisions p.a.* (and, indirectly, collision probability) was performed as early as 1987 in the form of conversion factors (Hydén, 1987). However, these conversion factors are still limited by the aforementioned constructed issues of TCT, and the small scope and the age of study, though this does suggest that satisfactory validation is possible.

In the meantime, TCTs have seen a resurgence in usage in the past decade, both in Northern Europe and in North America, with the introduction of analysis techniques through accessible and affordable computer vision (Saunier and Sayed, 2006b; Laureshyn et al., 2009) and modernization of the surrogate safety method (with Saunier et al., 2007; Tarko et al., 2009; Laureshyn et al., 2010; Laureshyn, 2010) and others.

2.3.2 Modern Use of Surrogate Safety Methods

An alternative school of thought regarding traffic conflicts (e.g. Amundsen and Hydén, 1977) gave rise to methods inherited by the modern surrogate safety methods. Modern surrogate safety methods retake the original idea of traffic conflicts as precursors to collisions, however they use measures designed to describe objectively defined road user interaction intended for use as safety indicators, i.e. a collision precursor. These measures include, notably, time-to-collision (TTC) (first proposed by (Hayward, 1971)), post-encroachment time (PET), Required Deceleration, Deceleration to Safety Time (DST) (Hupfer, 1997), Proportion of Stopping Distance, Deceleration Rate to Avoid Crash, and many others, as potential safety indicators for collision probability.

- **time-to-collision (TTC)** is a very popular SSM which measures the remaining time before a collision course results in a collision. It is based on the principle that, as the TTC measured for a particular collision course decreases, the chance that at least one road user reacts to or engages in evasive manoeuvres before a collision ensues decreases dramatically (Hayward, 1971), since the reaction time of individual humans is static (with the exception of such effects as chances of perception). This critical reaction time is commonly measured at 1.5 s on average in the literature (Hydén, 1987; Green, 2000), although reaction time varies substantially throughout the human population and under different conditions. TTC is measured in a number of different ways, depending on how exactly a collision course is defined in the literature. The most basic application is as follows:

$$TTC = \frac{\|\vec{d}\|}{\|\Delta\vec{v}\|} \quad (2.3)$$

where \vec{d} is the distance vector separating two road users and $\Delta\vec{v}$ is the differential velocity vector between the two road users. Note that this equation only applies if the collision course motion is presumed to be constant and linear. New, more sophisticated calculation methods and assumptions of motion in a collision course will be presented later as part of the contributions of this work.

- **PET** is the time difference between successive arrivals of overlapping real observed trajectories at the point of overlap (Allen et al., 1978). This measures acceptable encroachment road user behaviour analogous to traffic gaps accepted during gap acceptance (also measured in units of time).
- **predicted post-encroachment time (pPET)** is measured similarly to PET, only continuously, using trajectories in collision courses (the same ones that are used to calculate TTC) instead of observed trajectories. A special, related form of pPET is sometimes also referred to as **gap time (GT)** and follows the definition: “The time lapse between the completion of encroachment by a turning vehicle and the arrival time of a crossing vehicle if they continue with the same speed and path.” (American Association of State Highway and Transportation Officials, 2010).
- **Proportion of Stopping Distance** represents the “ratio of distance available to manoeuvre to the distance remaining to the projected location of collision” (Gettman and Head, 2003). It is less frequently used.
- **Deceleration Rate to Avoid Crash (DRAC)** (Guido et al., 2011) is defined as “the

minimum deceleration rate required by the following vehicle to avoid a crash” (Kuang et al., 2015), i.e. it is the ratio of the differential velocity to the TTC. A very similar concept, if not related, has been previously proposed under a slightly different name as “deceleration power needed to avoid an accident” (as in Hydén, 1987). “Deceleration Rate” is sometimes also mentioned in the literature (e.g. Gettman and Head, 2003), though its context is not always as explicitly stated. Substituting equation 2.3, DRAC becomes

$$DRAC = \frac{\|\Delta\vec{v}\|^2}{\|\vec{d}\|} = \frac{\|\Delta\vec{v}\|}{TTC} \quad (2.4)$$

The typical DRAC threshold of safety ranges from 3.35 to 3.4 m/s^2 (AASHTO, 2004; Archer, 2004).

Similarly, measures of speed have been used as safety indicators for collision severity. These measures have been recently re-explored using automated techniques (in Saunier et al., 2007; Saunier and Sayed, 2008; Ismail et al., 2010; St-Aubin et al., 2014).

Though the exact link between these various safety indicators and collision probability is, for the time being, still largely untested, it is generally accepted in the literature that the most popular safety indicators have at least some intuitive relationship with collisions. This might make them useful as qualitative predictors of collision probability today and potentially quantitative predictors in the future. Simple conversion factors have been calibrated for traffic conflict-to-collision ratios linking collision rate with traffic conflict rate (linearly) in a sample of Swedish TCT (Hydén, 1987).

So far, surrogate safety methods have, subject to minor road design modifications, largely been used in limited capacity in before/after studies of controlled sites, for example: slip lanes (Autey et al., 2012), scramble-phase movements at an intersection (Ismail et al., 2009), and special highway lane signalization treatments (St-Aubin, 2012). Despite this, challenges exist when designing studies using these measures such that indicators remain comparable across the subject of research and control sites.

The Surrogate Safety Assessment Model (SSAM) was an initiative to compute various safety indicators from microsimulation software (Gettman et al., 2008). However, SSAM makes use of simulated vehicle trajectories and is somewhat controversial. Surrogate safety methods are ultimately the study of outlying road user behaviour. SSAM is basically incompatible with the modernized surrogate safety methods framework proposed as part of this work because

- microsimulation software uses “perfect driver models,” i.e. models that cannot cause road users to collide with one another, that have perfect knowledge of their positions

and speeds and that of the surrounding road users, and do not reproduce anomalous road user behaviour;

- in practice, practitioners rarely calibrate driving behaviour parameters properly (see Rakha et al., 1996) and to the degree that would be required to reproduce anomalous road user behaviour; and
- validation results found for the model were low to moderate at best, with a collision to traffic conflict correlation of 0.41 (Gettman et al., 2008).

More recently, several general evaluation frameworks have been proposed, including probabilistic road safety analysis frameworks (Saunier and Sayed, 2008), and the systematic redefinition of the traffic conflict concept in the form of the modernized safety indicators derived from the basic concept of a collision course which makes use of motion prediction methods (Mohamed and Saunier, 2013; Mohamed, 2015). The collision course, which follows from the original idea behind the traffic conflict, is defined as an instantaneous traffic event involving exactly two road users (or, theoretically, one road user and one non-moving object) whose motion-predicted trajectories collide in the very near future.³ Under the most generic definition, the collision course is the set of colliding trajectories for each road user involved at an instant in time (interaction instant). Since surrogate safety methods are substituted for the observation of real collisions, these colliding trajectories are modelled using motion prediction at the interaction instant. This, in turn, is the process of modelling expected road user behaviour, which process leads to the motion of the road user where dynamic evasive action does not take place. For example, under the same collision course conditions, what is the outcome if a road user is momentarily distracted? Work by Davis *et al.* suggests that evasive action undertaken during a collision course is different from that undertaken during a collision (Davis et al., 2008).

2.3.3 Collision-Course Motion Prediction

The earliest application of collision course motion prediction (although conceptually and implicitly understood as such at the time, the definition has expanded considerably) was constant velocity motion prediction. This method of motion prediction proposes that collision courses are formed when road users travel “with movement remaining unchanged” (Amundsen and Hydén, 1977), which is interpreted today as meaning “speed and direction remain unchanged,” since a non-zero displacement of road users is still a precondition for

³In the context of road safety, *near future* implies time scales close to that of human reaction, decision, and execution times, generally on the order of magnitude of 10^0 to 10^1 seconds.

collision. The basic idea behind this premise is that, ignoring minor speed losses, road users obey Newton's First Law of Motion unless control input, i.e. control input in the form of evasive action taken in a collision course, is received.

However, this type of motion prediction makes a strong and restrictive assumption about the natural motion of road users and the conscious control input of road users, and is not deemed suitable for all road environments, notably non-linear environments such as roundabouts. For example, Newton's First Law of Motion (constant velocity) is not in effect when a motor vehicle is engaged in a turn, since pavement friction applies rotation to the vehicle if any wheels are tilted (as is the case of a motor vehicle engaged in a turn). Newton's First Law of Motion does not apply to pedestrians in the same way, but it is reasonable to assume that pedestrians will continue to walk in a general direction if they are not actively avoiding an obstacle (e.g. not seeing an oncoming car while crossing in a cross walk).

A large body of motion prediction and collision avoidance work has been conducted in the field of robotics in two major model families: deterministic, where motion prediction is represented as a single future trajectory, and stochastic, where motion prediction is probabilistic in nature, among a set of possible future trajectories (Eidehall and Petersson, 2008).

Saunier proposed (in Saunier et al., 2007; Saunier and Sayed, 2008; Saunier et al., 2010) a need for collision courses to incorporate multiple probabilistic paths, giving room for more sophisticated motion prediction methods to arise. Recent work in collision course motion prediction specific to surrogate safety methods has been undertaken by Mohammed (Mohamed and Saunier, 2013) and a variety of context-free models has been proposed. Context-free motion prediction is performed exclusively from parameters intrinsic to the road user and ignores road conditions. The proposed methods include, in addition to constant velocity

- **Normal adaptation:** A motion prediction method that predicts motion with random small variations, probably unconscious, of a road user intending to move in a straight line overall (Mohamed and Saunier, 2013). This behaviour can be assumed to be either conscious or not.
- **Point Set Prediction:** This is a special application of other motion prediction methods, where each position of the characteristic tracking features of each road user are used instead of the aggregated position, thereby minimizing micro-corrections caused by a continually changing number of tracking features (Mohamed and Saunier, 2013). This is an implicit, rather than explicit, use of a road user outline and it only applies to trajectory data extracted using a feature-based motion tracker. This method is also known as **Set of Initial Positions** (Mohamed, 2015). Alternatively, the method could

be adapted to samples of shape with other tracking methods.

- **Evasive Action:** This method of prediction is a stress test of a driver’s ability to evade a collision in order to compute a new safety indicator, complementary to TTC: the probability of unsuccessful evasive action. This safety indicator samples the space of possible evasive actions (Mohamed and Saunier, 2013).

Some previous applications of surrogate safety methods assumed constant velocity trajectories on straight highway segments (St-Aubin, 2012), an arguably acceptable environment to use for constant velocity motion prediction, though such assumptions are limiting in non-linear driving environments. If trajectory data is available to train a motion prediction algorithm, motion patterns which model trajectory motion from learned motion probabilities of similar road users can be used instead. For example, empirical models learned from trajectory history have been developed, including expectation maximization of a hidden Markov model (HMM) for pedestrian classification and movement prediction (Bennewitz et al., 2005a).

Saunier developed trajectory prototypes and path prediction probabilities using longest common subsequence (LCSS) similarity in a probabilistic framework, and extended the TTC indicator by computing its expected value (in Saunier et al., 2007). These trajectory prototypes were further expanded (in Mohamed and Saunier, 2015) adjusting trajectory speed using learned speed profiles. Morris and Trivedi published a domain-independent comprehensive analysis of various trajectory learning algorithms (Morris and Trivedi, 2008), including HMM, K-Means, fuzzy C-means (FCM), Similarity Threshold, hierarchical methods (Agglomerative, Divisive), neural networks (SOM, Fuzzy SOM, SOFM) and co-occurrence methods (e.g. Document Keyword). Some of these approaches are used for motion prediction or classification, or both, by various researchers.

Hu implemented a Fuzzy neural network for motion prediction (Hu et al., 2004). Laugier used HMMs and Gaussian processes to estimate collision probability (Laugier et al., 2011). Implementation of these methods specifically for the purpose of collision course motion prediction remains limited. A much greater range of motion prediction has been explored in the field of robotics (notably Bennewitz et al., 2005b; Althoff et al., 2008; Sorstedt et al., 2011; Sekiyama et al., 2011).

These methods, as well as the new methods introduced as part of this research project are discussed in further detail in sections 4.3.1 and 4.3.2 of Chapter 4.

2.4 Roundabouts

In this work, roundabout safety and road user behaviour are studied for the purposes of the application of the new surrogate safety method, and as contributions in and of themselves. Roundabouts serve as an example of a type of road infrastructure where traffic interactions are managed entirely by road user behaviour (Rice, 2010). In fact, road user-managed stopping and yielding are a distinguishing feature of the roundabout, especially in North America, where unsignalized intersections are exceptionally rare. The implementation of roundabouts in North America is a relatively recent phenomenon, unlike in Europe. For example, roundabouts in the province of Québec have existed only since 1998. Meanwhile, the first roundabout built in the United Kingdom, where the modern form of the roundabout was codified by the United Kingdom’s Transport Research Laboratory, dates back to 1966.

It is important to distinguish between roundabouts and traffic circles (sometimes also called *rotaries*). The former are generally more compact and unsignalized, and require entering vehicles to yield at (nearly) orthogonal approaches, while the latter are generally much larger, typically includes signalization at each approach, and do not provide requirements for entering vehicles to yield, as the approach is tangential and signalized. These differences produce important differences in road user behaviour (Rice, 2010), as interactions in roundabouts need to be managed by individual road users, whereas in traffic circles, interactions are managed by explicit attribution of the right-of-way through signalization control.

Roundabouts are typically sold on alleged merits of road safety. This will be covered in detail in the next section. However, in addition to matters of road safety, other operational benefits are frequently cited for roundabouts (Rodegerdts et al., 2010). Roundabouts

- create fewer delays;
- have better integration with existing traffic light coordination and thus require fewer queuing lanes; and
- lack the complexity and maintenance required by traffic lights;

but also have some disadvantages (Rodegerdts et al., 2010):

- for the same volume of traffic, besides small reductions in queueing lanes, roundabouts require a significantly larger footprint;
- roundabouts have a maximum theoretical per-lane capacity smaller than that of ordinary traffic lights in most situations; and

- multi-lane roundabouts have a number of safety and performance issues, and thus tend to scale poorly with an increasing number of lanes (though overall, still feature a comparable or better safety record).

Given the small number of roundabouts, their relatively recent deployment in North America, and an overall different design philosophy with respect to intersection signalization in North America, issues with driver culture, driver education, and general road safety have been raised with roundabout adoption in North America. In response, several design guides have been published on the subject, including the MTQ's guide (Ministère des Transports du Québec, 2002) and the Transportation Research Board (TRB)'s sponsored National Cooperative Highway Research Program (NCHRP) Report 572 (Rodegerdts et al., 2007) and NCHRP Report 672 (Rodegerdts et al., 2010). However, the number of North-America-specific studies remains small, especially in regards to understanding the underlying mechanisms of road safety.

2.4.1 Roundabout Safety Studies

Roundabouts are typically sold on alleged merits of safety. In addition to reducing speed, and thus collision severity, in Swedish studies (Hydén and Várhelyi, 2000), they have been proposed as a method for managing conflict points at intersections. Generally speaking, the design principle of the roundabout is to provide at-grade intersections with fewer points of conflict,⁴ as illustrated in the commonly reprinted Figure 2.1 (Rodegerdts et al., 2010). However, this type of analysis is questionable, and possibly even deceiving: while fewer *points* and *types* of traffic conflicts are depicted in Figure 2.1, it does not necessarily follow that fewer traffic conflict *events*, and thereby fewer collisions, are occurring for the same volume of traffic.

In fact, arguably more traffic conflicts are produced in the form of merging manoeuvres, because road users must yield (interact) with one another, where as, at a traffic light, interactions between road users are regulated. Stated more simply, the expected, and predictable, behaviour of road users at a traffic light is to stop, and right-of-way is never contested, whereas at the yield line of a roundabout, it is less obvious which road users have the right of way. Thus the following question arises: is it better to generate a small number of diverse types of traffic conflicts or a large number of similar traffic conflicts? This question can be pushed further by substituting *type diversity* from the previous statement with *seriousness*.

⁴Note that the *point of conflict* concept is distinct from the *traffic conflict* concept presented earlier. Whereas *traffic conflict* are observed traffic events, points of conflict are locations, where particular types of interaction are theorized to exist.

Such questions have yet to be investigated further in the surrogate safety methods space, and road safety space more generally.




Figure 2.1 Conflict point comparison between a regular intersection and a roundabout (c) (Rodegerdts et al., 2010).

Accident Report-Based Studies

Despite reducing the number of conflict points (Rodegerdts et al., 2010), there is still some debate regarding the practical reduction in collisions, both at the global and at the local level. However, recent studies based on historical accident analysis, showing reductions in vehicular collision severity, suggest that vulnerable road users are still at considerable risk (including Hydén and Várhelyi, 2000; Persaud et al., 2001; Gross et al., 2013), and a study by Jensen found a decrease in motor vehicle collisions but an increase in cyclist collisions (Jensen, 2013). A number of roundabout collisions studies have been compiled (see Rodegerdts et al., 2007, 2010).

Daniels *et al.* found that accident rates differed from roundabout to roundabout for vulnerable users (pedestrians and cyclists) and were highly correlated with traffic exposure in Flanders, Belgium (Daniels et al., 2010). Similar figures have been demonstrated in Victoria, Australia (Cumming, 2012).

It is interesting to note that roundabout centre island height has been cited as improving safety (Jensen, 2014). This result is counter to classic safety models. Jensen argues that

sight distances at the merging zone are sufficient, and that the view of road users on the other side of the roundabout is superfluous information and, moreover, acts as a distraction. Alternatively, if the view distances are not sufficient, he argues that road users act more cautiously as a result.

Chen *et al.* (in Chen et al., 2013) found that average approach speed was the most significant predictor of number of collisions, using Bayesian Poisson-gamma and zero-inflated Poisson models to predict collisions as a safety performance function. This study also found that roundabout diameter correlated with average approach speed, suggesting increases in expected collision probability (this inference is not entirely compatible with the safety framework presented earlier).

Surrogate Safety-Based Studies

Use of surrogate safety methods for roundabouts is more limited.

Hydén and Várhelyi studied 21 roundabouts and found that speeds always decreased four months after implementation of a roundabout, although some gains in speed reduction were lost after four years (Hydén and Várhelyi, 2000). Roundabout speeds tended to stabilize around 30 km/h; one roundabout approach with a previous operating speed of 20 km/h saw its operating speed increase after implementation of the roundabout, suggesting that roundabouts might have a fixed influence on speed in their environment. The study suggested that these changes in speed had a negative impact on travel time and emissions for major streets, with gains on smaller streets.

The Swedish TCT was used by Sakshaug *et al.* to study traffic conflicts of cyclists and motorists in roundabouts. It was concluded that cyclists within a roundabout conflicted with motorists exiting the roundabout if they were driving side-by-side, while cyclists crossing the approach and exit next to the crosswalk caused yielding ambiguity.

An analysis of observed gap acceptance (e.g. acceptance of GT, used mostly in car following models) in a saturated multi-lane roundabout in Lund was performed using some manually-annotated video data to calibrate microsimulation software (Irvenå and Randahl, 2010). Although gap acceptance will be revisited in this work in a derived form primarily as an SSM of yielding, gap acceptance also serves to calibrate flow models for basic traffic analysis purposes, as it is a general-purpose measure of traffic behaviour.

Al-Ghandour studied roundabout slip lanes using SSAM along with Poisson regression and concluded that slip lanes reduce traffic conflict occurrence (Al-Ghandour, 2011), though the question of whether this reduction in traffic conflicts leads to a reduction in collisions is not

thoroughly addressed. It should be noted that the NCHRP design guide classifies roundabout slip lanes as non-standard, since they can induce conflicts between cyclists and pedestrians (Rodegerdts et al., 2010).

Of particular note is a thesis recently written by H. Sadeq (Sadeq, 2013) who studied a single roundabout extensively using video tracking and TCT. The methodology still relied heavily on manual interaction interpretation and a form of serious event comparison (SEC). Use of SEC is problematic, as will be demonstrated in Chapter 4, but the work by Sadeq is still remarkable for exploring traffic violations at roundabouts in great detail.

The perceived safety of roundabouts is also mixed in the context of North America. Jacquemart and Pellecier and St-Jacques cite issues of driver education, practical complaints about the design, and uncertainty regarding the safety benefit (Jacquemart, 1998; Pellecier and St-Jacques, 2008). However, Jacquemart argued that opinions would change upon implementation as users had an incomplete or incorrect understanding of what a roundabout is and how it works. Retting *et al.* studied the effects of various factors on perceived safety, including signalization and design (Retting et al., 2007) and posited that the most prevalent problem was unsuitability of the design to the particular requirements of the intersection. Meanwhile, Perdomo *et al.* conducted stated preference surveys with vulnerable road users and concluded that the presence of pedestrian-oriented design features, such as cross walks and crossing lights, had a large effect on pedestrian preference (Perdomo et al., 2014), though it remains to be seen if this effect is large enough to effect a significant modal shift or trip re-routing.

Summary of Literature Shortcomings

Overall, there are no shortages of surrogate safety method, however issues of systematic definition of the concepts and measures abound. Furthermore, most existing motion prediction methods are not apt at modelling road user behaviour in complex, non-linear environments, such as roundabouts.

Progress has been made in the area of computer vision for purposes of traffic analysis, but none of a size and on a scale that might be required by a complete surrogate safety method framework. Such a system needs to be able to collect high-resolution traffic data for all road users in a system passively, and it must scale efficiently to deployment at numerous sites. Scaling of these high-resolution systems is typically bottle-necked by equipment installation and a need for efficient general-purpose data validation. These challenges are to be addressed with the system presented in the next chapter.

As for roundabouts, there is still much to learn, specifically regarding the microscopic effects of design on road user behaviour and the suitability of design to North American culture (e.g. Retting et al., 2007). No research has ever attempted an international comparison of road user behaviour at the proposed level of detail. Some macroscopic road safety studies performed in Europe may not apply, at least in the short term, if local variables and road user culture or attitudes can explain some degree of difference in road user behaviour and road safety.

CHAPTER 3 A NOVEL FRAMEWORK FOR LARGE-SCALE HIGH-RESOLUTION TRAFFIC DATA COLLECTION USING COMPUTER VISION

This chapter expands on section 2.1.1 to propose a framework for a complete, automated, video-based traffic data collection and analysis system. While numerous traffic data collection systems have been devised using a wide array of technologies, few currently existing traffic data collection systems are capable of recording high resolution trajectories of *all* road users at a given site, and those few systems that do (e.g. Gordon et al., 2012) are not adapted for projects requiring efficient large-scale data collection. The intent of the system presented in this chapter is to do both simultaneously: collect microscopic road user trajectory data using a system that scales efficiently for data collection at a large number of sites. Typically, equipment installation and data validation overhead are limiting factors in scaling microscopic data collection systems to applications that require observations collected at numerous sites. The chief motivating goal behind the development of this system of high-resolution trajectory data collection lies in the need to efficiently capture data at numerous sites for the purpose of road user behaviour analysis, road safety studies, and the surrogate safety method presented in the next chapter.

This system collects the trajectories of all road users passively, regardless of road user type (motor vehicles, trucks, cyclists, pedestrians, etc.), and anonymously. Each trajectory collected by the system measures the spatial positions of a road users to a high degree of accuracy, virtually continually (typically more than 15 positions per second). To achieve this goal, trajectories are automatically extracted from video data using computer vision techniques. The intent of this chapter is first, to list the design goals of such an automated video-based traffic data collection system, and then, to guide the reader through the basic methodological process of collecting and processing large quantities of video data, and finally, to equip practitioners with practice-ready tools for the large-scale analysis of traffic data. Since the focus is on large-scale, detailed data collection, particular attention is brought to issues related to performance, data storage, and data management. The data collected and prepared following these procedures forms the basis for the work presented in the remaining chapters.

This chapter is not intended to review computer vision algorithms or enhance computer vision theory, as the bulk of this work can be found elsewhere (Bradski, 2000; Jackson et al., 2013). Furthermore, readers are not expected to be proficient in the technical aspects of computer vision theory. Many of these tools are readily available. In fact this chapter builds

on top of the open-source *Traffic Intelligence* project (Saunier and Sayed, 2006a; Jackson et al., 2013) which already implements a fully-featured trajectory extraction algorithm and a trajectory data storage scheme. Instead, this chapter is devoted to issues related to applying these computer vision tools to large volumes of video data for conducting specialized traffic engineering studies.

3.1 System Overview

The objective of a fully automated, large-scale video-based traffic data collection system is the collection of high resolution trajectory data at many different sites for use in traffic and road safety analysis applications. Figure 3.1 provides a general roadmap for this process. The basic video-based traffic data collection system is designed to be integrated in an ordinary traffic study or road safety analysis, typically initiated with an inventory of sites to be studied, followed by a selection of sites to sample following ordinary choices governing statistical sampling in any type of study. Once the site selection has been completed, the traffic data collection activities can begin. To achieve automation, not only is computer vision used for automating the trajectory extraction process, but the entire set of video data and corresponding video data is stored and indexed. The only manual operation of the system is the recording of video data, if traffic cameras are not already available in place, tracking calibration as necessary, and metadata entry, if this is not readily available.

The data collection system collects and processes traffic data in four distinct stages. To begin the process, video data is collected at each site either through existing traffic cameras or using a purpose-built video data collection unit (VCU). At this stage of the system, any issues with the video data, such as lens distortion, vibration, and encoding issues, are corrected. Along with camera parameters and scene metadata, the process of extracting road users trajectories is completed using computer vision algorithms, i.e. *feature-based tracking* (Bradski, 2000; Jackson et al., 2013). At this stage, only **moving objects** are identified, each represented by a trajectory within the image. It is only once data clean-up and validation are performed that these moving objects are referred to as **road users**, forming the bulk of the study data.

These trajectories provide a quasi-continuous series of the position of each road user travelling through the scene being studied. Although position observations are discrete, they replicate trajectory continuity through a very high sampling rate. These positions are projected onto a Cartesian plane corresponding to a map of the site. Since the position data is quasi-continuous, derived data such as quasi-instantaneous speed, heading, and acceleration can be obtained. However, it should be noted that, at this time, acceleration data collected

with current tracking mechanisms tends to be very unreliable without heavy filtering (e.g. as found in Punzo et al., 2009). More importantly, successive errors add up at each derivation between frames, but may be corrected with a Savitzky-Golay filter. As tracking quality improves, acceleration accuracy is expected to improve to a usable state without filtering.

Data interpretation is fully automated and integrated with the data collection system. However, the specifics of this process are covered in later chapters, as this will require a special discussion of road safety models and geometry-specific parameters (Chapter 4 and parts of 5 respectively).

The process of feature-based tracking is entirely automated as well. The degree of automation of the entire system thus largely depends on the degree of curation of input data, including scene metadata and camera parameters; automation is also improved with more accurate feature tracking, as less intervention is necessary to correct errors.




Figure 3.1 Overview of the data collection pipeline embedded in a typical traffic engineering study using video data.

3.2 Video Data Collection

The bulk of the data collected as part of the data collection system consists of video data. At this stage, the quality of the source video data is of concern. While a significant amount of error correction can be made via software, much trouble can be saved and data accuracy improved by considering the following video data requirements:

- The video must have a constant framerate. This is the foundation for nearly all implementation details of moving objects as these use frame index to track the dimension of time. More importantly, however, is the challenge that a variable framerate poses when interpreting measures defined instantaneously, and “continuously” over a series of instants. This will be important in the discussion of the next chapter. Cameras and codecs with variable or dynamic framerate are therefore challenging to work with. Alternatively, variable framerate can occur if the camera overheats or if video is encoded at a bitrate nearing or exceeding the limits of the camera’s processor and frames are skipped. Many consumer-grade cameras have these issues, as consumers tend to favour image quality over imperceptible—to the human eye, at least—issues with framerate.
- The video requires an adequate resolution such that enough identifying features can be found of road users captured on video by the tracking algorithm. However, higher resolutions provide diminishing returns on tracking accuracy and tracking distance, and are more computationally expensive to process, store, and handle.
- Since each video frame is used as an observation of time, the video data must have a framerate high enough to qualify the extracted trajectories as quasi-continuous. This generally means that no displacements between successive observations should be greater than the size of the moving objects being tracked. Modern consumer-grade cameras record with framerates as low as 15 frames per second (FPS) but can also record as many as 60 FPS or more. As with video resolution, more FPSs lead to increased tracking accuracy—more specifically, trajectories replicate continuity with greater fidelity as more discrete points are sampled per second—at the cost of increased storage requirements and processing time. At 15 FPS, a moving object travelling at 30 km/h will have an inter-sample spatial displacement of $0.5\bar{5}$ metres per frame. As a general rule of thumb, inter-sample displacements larger than the size of the road user should be avoided because these larger inter-sample displacements could potentially cause issues with trajectory-based data interpretation.
- The video data must be taken from an elevation sufficient enough to minimize parallax

tracking error and vehicle occlusions. This and other practical issues with weather and tamper-proofing are discussed further in the next section (3.2.1).

Note that the video data need not use only visible light. While the concepts and experimental results presented in this chapter and remainder of the work are built around the use of ordinary visible-light cameras, alternative forms of optics may be used as well, provided that enough contrast is captured between the background and road users. Example applications of alternative optics include thermal cameras for low-light or poor visibility conditions. These have been briefly tested with the system (Fu et al., 2016).

The video data is typically complemented by a suite of metadata. First, the frame rate of the camera must be known and recorded. Optionally, lens distortion parameters can be used to correct any distortions in the recorded scene caused by a wide-angle lens (see section 3.2.2). Satellite images of the scene where video data is collected, while not strictly necessary, provide useful geometric and geographic information through the use of projection transformation (see section 3.3.2). Expansion factors can be optionally imported to automatically calculate average annual daily traffic (AADT) during traffic analysis. Sequence start times should be known, especially when undertaking time of day analysis, including AADT, or where multiple cameras at the same site must be synchronized. Finally, any necessary engineering factors to be evaluated can be cross-referenced for the purpose of a road safety analysis.

3.2.1 Equipment

Many iterations of the video data collection unit (VCU) (Jackson et al., 2013) were built and tested during this project. The first of these was built during an earlier iteration of this system for a project with similar requirements (St-Aubin, 2012). The VCU is a telescopic pole that attaches to a fixed street pole. A fixed weather-insulated camera head is placed at the top of the pole, and power and Ethernet are passed up through the pole from the base station where a video feed can be previewed during VCU installation. This VCU is then secured to the street pole to deter theft and tampering, and to stabilize the camera. This design underwent many iterations to address a host of design challenges and requirements. In addition to meeting the data specifications of section 3.2, which are constrained by power, heating, and weight considerations, the VCU must be

- discreet, so as to not distract or frighten road users;
- independent so that it can be installed in any remote environment, i.e. having its own

data storage and power supply capable of lasting at least 24 hours, or until the battery and data storage solution can be replaced;

- elevated in order to minimize parallax issues and issues with visual obstructions;
- portable, for easy and timely installation, with minimal equipment and no elevated work platforms;
- stable, so that vibrations from wind and traffic are minimized;
- weather proof and operable in both freezing and hot conditions; and
- tamper- and vandalism-proof.

Figure 3.2a demonstrates the generic installation of the VCU and Figure 3.2b demonstrates the installation of multiple cameras at the approach of a roundabout for data collection in Chapter 5. The pole is secured to the mast with rope, to eliminate vibrations and sway in the wind, and with chain, to deter vandalism. This configuration is very discreet from the viewpoint of drivers (the camera head is too high to be seen) and mostly ignored by pedestrians in busy areas.

The latest iterations of the VCUs were built with lighter fibreglass materials and taller, more rigid telescoping poles. Furthermore, Internet-Protocol-enabled (IP) cameras were replaced with GoPro cameras with increased on-board storage, video quality, and integrated Wi-Fi video previewing features, simplifying installation dramatically (Jackson et al., 2013). Future improvements to the VCU may leverage advances in battery density, eliminating heavy batteries at the base of the unit, and units with IP cameras may be replaced with thermal cameras to improve low-light road user tracking (e.g. Fu et al., 2016). Drones have also been proposed as an alternative data collection system (Puri, 2005; Coifman et al., 2006) and are slowly gaining traction (Kanistras et al., 2013).

For Chapter 6, a special VCU was borrowed from a collaborating Swedish research team and deployed for data collection. This VCU was built on the chassis of a trailer and incorporated a hoist capable of lifting two IP Pan-Tilt-Zoom (PTZ) dome cameras at an elevation well over 15 m. Although much less discreet, and less portable than the VCU used for data collection for Chapter 5, this unit had the benefit of being capable of operating for long periods of time, eliminating most issues with stability, elevation, and field of view warping. This unit is demonstrated in Figure 6.2.



(a) VCU deployed at reduced height near an intersection (b) Two VCUs deployed at full height near a roundabout, for full coverage.

Figure 3.2 Example video data collection unit installation.

3.2.2 Source Video Data Treatment

Before video data can be used for trajectory extraction, the source footage must be treated for imperfections and varying qualities and formats of video, then stored and indexed. This process starts with a series of tests on a sample sequence to determine that frame rate is truly stable, that frames are not skipped, that the image is stable, that no sequences are corrupted, and that the video codec and video container formats are compatible with the software running the video-based traffic data collection system. These problems are difficult to tackle after data has been collected. Thus, it is always preferable to address these issues with any new camera equipment before recording begins.

Sequence Length

Some cameras, especially IP cameras, split stored video into minute-long sequences organized by time of day, while other cameras, usually consumer-grade cameras, dump fixed-sized sequences on disk or fixed-length sequences of duration, if the bitrate is constant. For this system, sequence lengths of 20 minutes or 60 minutes are typically used. This is achieved

using lossless video concatenation, on the hour.¹

This choice is justified because shorter sequences result in split trajectories more often. For example: given a typical road user dwell time of 8 s, one-minute sequences result in nearly 30 % of trajectory disruptions by number (probability of being within the first 8 s or last 8 s of a 60 s video) or 15 % by average trajectory length. Meanwhile, longer sequences become difficult to handle in memory as their file size increases. A typical one hour sequence generates, at the time of writing, 4 GB video files using H.264 encoding. The video data and trajectory processing could be streamed as a technical workaround to memory limits, but the benefits of having sequences longer than one hour do not justify the additional cost of implementation of such a solution.

Lens Distortion Correction

A particular kind of source video data treatment is lens distortion correction. Virtually all consumer-grade camera lenses distort their image to a certain degree and this can cause issues when determining world space coordinates of tracked moving objects. For best extraction results, the image must be as undistorted as possible. Distortion can be identified in any image by verifying if any line that is straight in the real world and that is recorded on video does not appear straight in the image. In the most extreme cases, distortion manifests itself as the “fish eye” effect.

To undistort an image, a geometric transformation is performed using a rectification map generated from a camera’s intrinsic camera parameters and distortion coefficients (Brahmbhatt, 2013)—these are static parameters associated with the camera, its lens, resolution, and aperture settings. The intrinsic camera parameters take the form of the matrix A

$$A = \begin{bmatrix} f_x & 0 & c_x \\ 0 & f_y & c_y \\ 0 & 0 & 1 \end{bmatrix} \quad (3.1)$$

where c_x and c_y are the coordinates of a principal point that is usually at the image centre, i.e. the half-width and half-height of the camera’s resolution, respectively, and f_x and f_y provide the focal length of the camera’s x and y image dimensions. The distortion coefficients take the form

$$distCoeffs = [k_1, k_2, p_1, p_2, k_3] \quad (3.2)$$

¹For this project, Mencoder is used.

where $p_1 = 0$ and $p_2 = 0$ in this application (these parameters only apply for stereo rectification maps) and k_1 , k_2 and k_3 are calibration parameters of the lens. These parameters are calibrated by first applying a reverse rectification map on sample footage of a checkerboard (the lines and corners of the checkerboard are easy to identify, and must all appear straight in the image). The rectification map maps undistorted coordinates (u, v) to source (distorted) coordinates (x, y) and takes the form, for each coordinate:

$$x = \text{map}(u, v) = u(1 + k_1r^2 + k_2r^4 + k_3r^6)f_x + c_x \quad (3.3)$$

$$y = \text{map}(u, v) = v(1 + k_1r^2 + k_2r^4 + k_3r^6)f_y + c_y \quad (3.4)$$

where r is the radial distance

$$r(u, v, c_x, c_y) = \sqrt{(u - c_x)^2 + (v - c_y)^2} \quad (3.5)$$

of the source point at (u, v) from the centre of the image, defined by c_x and c_y .

Once these parameters are obtained and the rectification map parametrized, it can be applied to any footage of the same camera with identical recording settings using an ordinary image remapping function, such as the remapping function included with OpenCV (Brahmbhatt, 2013).

An example of undistortion of a single frame can be seen in Figure 3.3 (with annotated points).



Figure 3.3 Side-by-side comparison of image before (left) and after (right) lens distortion correction.

Since the process of undistortion is systematic, it is applied to all cameras automatically.

Undistortion is implemented directly alongside the feature-based tracker implementation so long as, for each camera type, intrinsic camera parameters and distortion coefficients are determined and stored beforehand. This has to be done only once per VCU.

3.2.3 Video Data Storage and Indexing

Video data is quite large and requires particular attention to storage. For example, for the case study performed in Chapter 5, 1874.0 GB of video data is collected, and an additional 25 %, or 481.3 GB, of storage is needed for data analysis overhead. The size of data permits storage on a single consumer drive, optionally with copies on multiple computers to (i) act as redundancy and (ii) increase performance, since processing can be divided among systems. For a larger project, it might be conceivable to store a master copy on dedicated network attached storage hardware with off-site backup and to parallelize all tasks over the network. This would require server-grade networking hardware.

The video data files are stored following a specification which allows for the systematic and categorized access to files by computer program. They are then indexed with metadata in a portable database file (see sections B.2.1 and 3.4.1 for specific implementation details). The video files are separated into folders according to site and time of recording. Further technical details are provided in section B.1.

3.3 Vehicle Tracking

This section describes the objectives and methods of the trajectory extraction component of the traffic data collection system after video data has been obtained and indexed. This system combines (i) the computer vision algorithm, which performs the tracking of captured road user motion in the form of pixels within the video data and forms projected moving objects with real spatial coordinates, with (ii) validation of these moving objects such that they accurately reflect the real trajectories of the road users being studied.

3.3.1 Feature-Based Tracking

Two large families of tracking algorithms exist: feature-based tracking and background subtraction tracking (Buch et al., 2011). Ultimately, feature-based tracking, specifically the Kanade-Lucas-Tomasi feature-based tracker (Tomasi and Kanade, 1991), is used, as it is readily available in implemented form for traffic applications (Saunier and Sayed, 2006a; Jackson et al., 2013).

The Kanade-Lucas-Tomasi tracker identifies individual characteristic features of moving objects and tracks their motion throughout a continuous sequence of video frames, so long as the features do not remain stationary. The features are grouped together by similarity of motion forming moving objects representing the real motion of a road user. The output of this process can be seen in Figure 3.4. The red line, i.e. the **trajectory** proper, represents the average tracked motion of the feature positions corresponding to the recorded road user being tracked. The features are not shown in this figure, as they would obscure the car, but their location in the current frame is bounded by a blue square for each moving object.

The benefit of using the grouped feature trajectory over each feature trajectory individually is that the centre position of the moving object is implicit. The feature-based tracker does not know what part of a road user an individual feature is representing. Even worse, the relative location of individual features can change if the object rotates. However, one problem with grouped feature trajectories is that the number of detected features is not constant from one frame to another, and the centroid of these feature positions can jump from position to position. In this traffic data collection system, grouped, rather than individual, feature trajectories are used for most operations, due to their intrinsic benefits and because the limitation related to variable length features is only bounded by the quality of the data (individual features disappearing or appearing along the trajectory have a less pronounced effect if, overall, there are more features to dilute the error).

In the case of Figure 3.4, tracking occurs for two distinct road users: a dark grey SUV and a white sedan. Note that the white van and the car behind it are not tracked, as they are too far from the camera to be accurately tracked.



Figure 3.4 Example feature tracking (after feature grouping) of two vehicles in the scene.

Larger video resolutions offer better tracking accuracy and maximum tracking distance, as

each moving object is represented on-screen by more visually distinct pixels and thus features. However, larger videos are exponentially more expensive to store and process. Experimentally, it is found that resolutions below 1024 by 768 pixels tend to be especially problematic. Resolutions at or just above 1024 by 768 pixels offer a decent compromise on quality and use of computational resources for zones of interest, where tracking is done, no further than 50 m away from the camera, though resolutions closer to 1920 by 1080 pixels may be preferable for tracking regions up to 100 metres away from the camera.

One last stage remains, and that is to transform the coordinates of all positions so that they lie in an overhead perspective, similar to those illustrated in Figure 3.10. At this stage,² the coordinates of the trajectories correspond to locations within the video frame (e.g. Figure 3.4). These will have to be transformed to world space coordinates next.

3.3.2 Data Structures

The following section documents how the trajectory data is structured and accessed. This is largely based on the methods presented by Jackson *et al.* ((Jackson et al., 2013)), however, it includes some additional updates and data types, including derived curvilinear positions.

Image Space

Decompressed video data is generally represented as a n_f by n_c by m by n matrix (array) containing integer values between 0 and 255. f is a sequence of frames and n_f is the length of that sequence. Given a known and constant frame rate, the sequence of frames effectively encodes time. c is a sequence of n_c channels encoding colour luminosity data. The most common colour encoding scheme by far uses three channels: red, green, blue (RGB).

An example of these different colour maps is illustrated in Fig 3.5. Each channel is displayed as a grayscale image, since individual channels contain luminosity data only. In RGB 8-bit colour graphics, colour intensity is encoded for each of the three colour channels with a value between 0 and 255. For example, 0, 0, 0 represents black, white is represented by 255, 255, 255, and 127, 0, 0 represents a deep burgundy red. m and n are, respectively, the height and width (resolution) of the video in pixels. Each pixel represents, in Cartesian space, the projected view of the world onto the camera plane in two dimensions. This space is called **image space**.

Image space has important applications for video annotation, but otherwise is of little use for traffic analysis. Note that coordinates of pixels in image space are discrete, due to the

²Note that, in practice, features are transformed to world space before grouping.

discrete nature of the video sensors used to capture images and the monitor technology used to display ordinary images. However, image space coordinates are continuous for everything else (e.g. annotations, tracked features, etc.)



Figure 3.5 Close-up of image space for each of the three RGB channels. Note how the red and green channels have much more contrast than the blue channel. The bottom triangle is the end of a yellow crosswalk, the top triangle is a curb, and everything else is pavement.

Tracked Objects

The bulk of the computer vision processes is devoted to extracting patterns in clusters of pixels in image space, and sometimes over successive frames, in order to organize this raw data into meaningful information, i.e. as moving objects. In the case of the feature-based tracker used for this project, the centroid of a moving object is tracked via characteristic features of that object over successive frames. A point p_{image} in continuous image space is represented by the corresponding Cartesian coordinates:

$$p_{image} = (x_{image}, y_{image}) \quad (3.6)$$

where x_{image} is the coordinate relative to the column index and y_{image} is the coordinate relative to the row index of the m by n matrix of pixels. Thus, a unit of 1.0 for x_{image} or y_{image} is equal to the width and height of a pixel respectively. Otherwise, these coordinates are decoupled from the matrix of pixels, meaning that a p_{image} can lie outside the matrix of pixels, and can lie within a fraction of a pixel.

While spatial data encoded from individual pixels is stored in the form of discrete values, features, feature groupings, and other computer vision data can in fact be represented as a fraction of a pixel, as demonstrated in Fig 3.5. Generally speaking, the fewer pixels representing an object, the greater the tracking error. Since the projected size of objects decreases

with distance from the camera (individual pixels represent effectively larger surfaces), spatial precision and tracking performance decreases as objects move further away from the camera and/or with low resolution video.

World Space

World space is the space corresponding to the real, unprojected three-dimensional positions of objects, measured in standard units of distance, i.e. meters. Since most terrain features, roads, and road users' movements all occur, for the most part, in two dimensions across a horizontal plane, the Cartesian coordinates of a point p representing one of these objects in world space can be simplified to the following two dimensions:

$$p = (x, y) \quad (3.7)$$

where x is an ordinary easting and y is an ordinary northing, both measured in standard units of distance (e.g. meters) from some arbitrary origin. The camera pinhole model is used to project from three-dimensional world space to two-dimensional image space, and vice versa, using the camera intrinsic and extrinsic matrices. For simple two-dimensional world space to two-dimensional image space coordinate transformation, a homography transformation is used (Jackson et al., 2013). The **homography** (and corresponding inverse homography) is a mathematical coordinate transformation which projects a point between two coordinate planes (e.g. image space and world space). This concept is illustrated in Figure 3.6 where the image space corresponds to the camera plane orthogonal to z_{cam} , and the world space corresponds to the red rectangle overlaid on the flat surface of the box, with example coordinates p_0 through p_3 . The 3-by-3 homography matrix \mathbf{H} thus relates p with p_{image} and is based on the following perspective transformation model for point i :

$$s_i \begin{bmatrix} x'_i \\ y'_i \\ 1 \end{bmatrix} = \mathbf{H} \begin{bmatrix} x_i \\ y_i \\ 1 \end{bmatrix} \quad (3.8)$$

between coordinates of image space (x_i, y_i) and world space planes (x'_i, y'_i) , where s_i is a fixed scaling factor. The homography \mathbf{H} is first built from existing planes using a back-projection minimization function, namely a screen capture from the video, and a satellite image scaled to standard units of distance in metres with matching static features between the two as reference points. The function to minimize is

$$\sum_{i=1}^n \left(x'_i - \frac{h_{11}x_i + h_{12}y_i + h_{13}}{h_{31}x_i + h_{32}y_i + h_{33}} \right)^2 + \left(y'_i - \frac{h_{21}x_i + h_{22}y_i + h_{23}}{h_{31}x_i + h_{32}y_i + h_{33}} \right)^2 \quad (3.9)$$

for a sequence of corresponding points $i = 0..n$. The implementation of the homography is readily available in the OpenCV library (Bradski, 2000).

In this way, coordinate (0,0) in world space normally corresponds to the origin coordinate (0,0) of the projection plane, i.e. the top left corner of the satellite image. If the geographical coordinate of the satellite image's origin is known, world space coordinates, effectively local coordinates, can be converted to geographic coordinates. In the case of Universal Transverse Mercator coordinate system (UTM) or Modified Transverse Mercator coordinate system (MTM), which uses a very similar system and distance of unit (metres), world space coordinates can be summed with the UTM or MTM coordinates to obtain geographic coordinates.

The bulk of traffic analysis uses world space coordinates. The remainder of the work will make use of world space coordinates exclusively, unless otherwise noted.




Figure 3.6 Illustration of a homography projection (c) Per Rosengren CC BY 3.0.

Curvilinear Coordinates

Curvilinear coordinates are an alternative to Cartesian coordinates in world space. Instead of using easting and northing, curvilinear coordinates make use of the orthogonally projected point p_{proj} on the nearest centre line of an individual lane *lane*. A point in curvilinear space

$p_{curvilinear}$ has the following three dimensions:

$$p_{curvilinear} = (lane, S, \gamma) \quad (3.10)$$

where S is the curvilinear distance to the projected point p_{proj} along the centre line of lane $lane$, and γ is the perpendicular distance between $p_{curvilinear}$ and p_{proj} (in a 2D plane, positive is to the right of the curvilinear vector). This is illustrated in Figure 3.7.

Curvilinear coordinates are useful for quickly determining generalized geometry-specific information, such as travel distance, travel lane, lane changes, and lane deflection. The basic algorithm for transformation f_{ct}

$$f_{ct} : (x, y) \Rightarrow (lane, S, \gamma) \quad (3.11)$$

from Cartesian world space coordinates (x, y) to curvilinear world space coordinates $(lane, S, \gamma)$ is provided in algorithm 1.

Algorithm 1 Curvilinear transformation

```

1: Input:  $p$ , alignments (a series of alignment objects, where an alignment is a series of
   segments, defined by an origin point and the origin of the next segment in the series).
2: Output:  $p_{curvilinear}$ 
3: function TRANSFORMTOCURVILINEAR( $p$ , alignments)
4:    $shortestDistance = \infty$ 
5:   for alignment in alignments do
6:      $S_0 = 0$ 
7:     for segment in alignment do
8:        $p_{proj} =$  orthogonal projection of  $p$  onto segment
9:        $\gamma =$  Euclidean distance between  $p$  and  $p_{proj}$ 
10:      if  $\gamma < shortestDistance$  then
11:        if  $p$  is to the right of segment then
12:          orientation = 1
13:        else
14:          orientation = -1
15:           $S = S_0 +$  Euclidean distance between  $p_{proj}$  and the segment's origin
16:           $p_{curvilinear} =$  ( alignment ID,  $S$ , orientation  $\times \gamma$ )
17:           $shortestDistance = \gamma$ 
18:        Increment  $S_0$  by segment's length
19:   if  $shortestDistance == \infty$  then  $p_{curvilinear} =$  None
   return  $p_{curvilinear}$ 

```

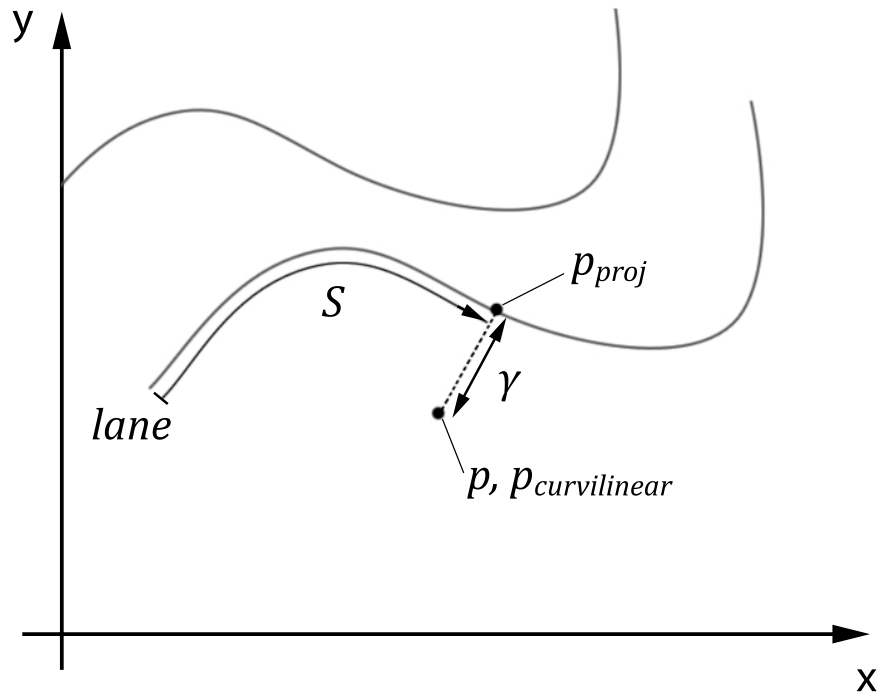


Figure 3.7 Comparison of Cartesian coordinates and curvilinear coordinates in world space.

For more information on lane centre lines and their segments, see section 3.4.2 on alignments.

Trajectories

A trajectory is simply an ordered sequence of n points through time t representing continuous motion between these points.

$$\text{trajectory} = (p_1, p_2, p_3, \dots, p_n) \quad (3.12)$$

Given a known and constant frame rate f and therefore a constant timestep Δt , i.e. $\Delta t = 1/f$, between the sequence of points, a point effectively has an implicit dimension of time t :

$$p = (x, y, t) \quad (3.13)$$

For example, for a video frame rate of 30 FPS, the time step between successive points is $1/30 \text{ FPS} = 0.033 \text{ s}$. Since time is known, and constant, the sequence of velocity vectors v_i of movement between successive points i and $i + 1$ can be easily derived from the differential of these successive points:

$$v_i = (v_{x_i}, v_{y_i}) = \left(\frac{(p_{x_{i+1}} - p_{x_i})}{\Delta t}, \frac{(p_{y_{i+1}} - p_{y_i})}{\Delta t} \right) \quad (3.14)$$

Moving Objects

In the automated video-based traffic data collection system, moving objects³ are the representation of road users' motion. Each moving object contains, among other things, a trajectory produced as a result of feature-extraction over successive frames of video data. The moving object's trajectory is stored in world space, and has derived curvilinear coordinates, a time series of derived velocity vectors, and start and end times of the sequence of points. The time between points is implicit according to the frame rate of the camera. The data models for positions, velocities, and moving objects (such as road users) as implemented in Traffic Intelligence as well as interaction models covered in the next chapter are presented in Figure 3.8 (St-Aubin et al., 2015a).




Figure 3.8 Moving object data model.

Note that in practice, the feature-based tracker used applies a moving average smoothing filter over the sequence of velocities of each feature. Also, for a trajectory of n frames and n positions, one effectively obtains $n - 1$ velocity vectors using this method. For the sake of simplicity, the first velocity vector is used to represent the first position and the last position always re-uses v_{n-1} as v_n . An offset of one frame represents no more than a fraction of a

³See section B.3.3 for technical details.

second, i.e. 0.033 s for a 30 FPS camera.

3.3.3 Validation

Before the moving objects can be used to represent road users for the subsequent chapters, trajectory data must be validated. In addition to imperfections in the current implementations of computer vision (e.g. the Kanade-Lucas-Tomasi feature-based tracker (Tomasi and Kanade, 1991)), the tracking process is prone to result in erroneous trajectories due to challenges in separating and grouping motion correctly, and problems with a number of different sources of optical interference, including variable lighting conditions, line-of-sight occlusion, and volumetric uncertainty projection (parallax) error. In general, many of these sources, but not all and not completely, can be mitigated by eliminating the source of interference and/or performing data collection under optimal conditions with appropriate conditions. For example, it is not feasible to eliminate sun or headlight glare during certain periods of the day when data collection is needed, e.g. at night and at dusk. To this end, two distinct methods of validation and error correction are performed on the trajectory data:

- Optimization of the feature-based tracker’s parameters is achieved by means of a specialized tracking accuracy test. This is a form of high-volume, low-impact quality control, processing a large number of trajectories, each undergoing minor correction. To evaluate tracking accuracy, the CLEAR MOT metric, Multiple Object Tracking Accuracy (MOTA) (and, to a lesser extent, Multiple Object Tracking Precision (MOTP), Bernardin and Stiefelhagen, 2008), are used to compare tracked objects with a sample of ground truth tracked objects obtained through manual annotation of the same video. MOTA is a robust, general-purpose, continuous, multiple object tracking accuracy metric. As such, it evaluates *all* instantaneous observed *and* missing points between tracked results and the ground truth. Optimization of tracking thus attempts to optimize this value via traditional optimization strategies for the input parameters of the feature-based tracker used. Results for the parameters of the feature-based tracker used are presented later in section 3.4.5.
- Once tracking has been optimized and undertaken, but before interpreting the data and making an analysis, error correction is performed in the form of a battery of filtering functions designed to target specific tracking errors. This is a form of low-volume, high-impact quality control as those few trajectories with the highest impacts on results are targeted for correction. These filtering functions are calibrated empirically and will be discussed later (section 3.4.6).

In addition to the aforementioned methods, quality of the data is controlled via manual inspection of random samples. Ultimately, the field of computer vision is constantly evolving. The feature-based tracker used as part of this video data collection system might be replaced with a more sophisticated tracking algorithm at a later date with better tracking performance, thus rendering these steps obsolete.

3.4 Semi-Supervised Model for Automated Traffic Analysis

For this project, Traffic Intelligence (Saunier and Sayed, 2006a; Jackson et al., 2013) is chosen as the computer vision platform on which to build the automated video-based traffic data collection system, itself based on the open-source computer vision OpenCV library (Brahmbhatt, 2013). The implementation of the collection of enhancements forming the automated video-based traffic data collection system, together with Traffic Intelligence, is known as *tvaLib*. It is both a library extending the functionality of Traffic Intelligence and a tool to assist and automate traffic analysis before, during, and after feature tracking takes place in Traffic Intelligence. This software implements a number of the features presented in this section, as well as the methodology presented in most of Chapter 4 and the case-specific implementation demonstrated in Chapter 5. However, a front-facing program implements typical module usage intended to perform common traffic analysis tasks. It supports up to 70 different commands, and the parameters of analysis, operation, and output can be configured separately. Additional functionality is included in a scripts folder to automate some post-analysis and other miscellaneous tasks. Technical information for operating *tvaLib* can be found in Annex B and attached with the software.

tvaLib adopts a semi-automated approach to traffic analysis. Although there is a fair amount of automation, there is also an important component of manual data input, particularly metadata annotation. Automation of some tasks is not justified, and in any case those tasks can benefit from human review. The library includes tools implemented for manual review of output and playback of trajectories so that targeted trajectories can be controlled for quality. This process is outlined in Figure 3.9. Some of these steps are potential candidates for automated implementation (for example, using trajectory clustering techniques or importing computer-assisted drawings to generate alignments), however the cost of implementation of a fully stable system is not currently justified.

One important contribution of *tvaLib* is the implementation of abstraction layers for metadata. The metadata is subdivided into five layers: sites, camera views, video sequences, site




Figure 3.9 Traffic analysis model demonstrating level of automation of the different tasks undertaken as part of the automated video-based traffic data collection system developed as part of this work.

analyses, and analyses. This abstraction is not only important for code clarity, it also organizes data collection hierarchy implicitly and traffic analysis explicitly. For example, a single site may be studied, but multiple cameras may be collecting video data disjointly: covering different portions of the scene, as demonstrated in Figure 3.10, at different times of the day. Furthermore, different types of traffic studies may be conducted at different locations within the same camera view using constructed **analysis zones**. The functionality of each abstraction layer is further documented in section 3.4.1. Note that tvaLib does not yet support *trajectory reconnection*. With this operation, trajectories from one camera view are merged with the trajectories from a different camera view at the same site, if they correspond to the same real object. However, tvaLib is built to support a future implementation of trajectory reconnection.




Figure 3.10 Sample vehicle trajectories from multiple camera views (2) at the same site. Each camera view is highlighted with a different colour.

3.4.1 Video Metadata

Video data is indexed in accordance with the following specification, itself based on and backwards compatible with an early metadata specification implemented in Traffic Intelligence and illustrated in Figure 3.11 (St-Aubin et al., 2015a).

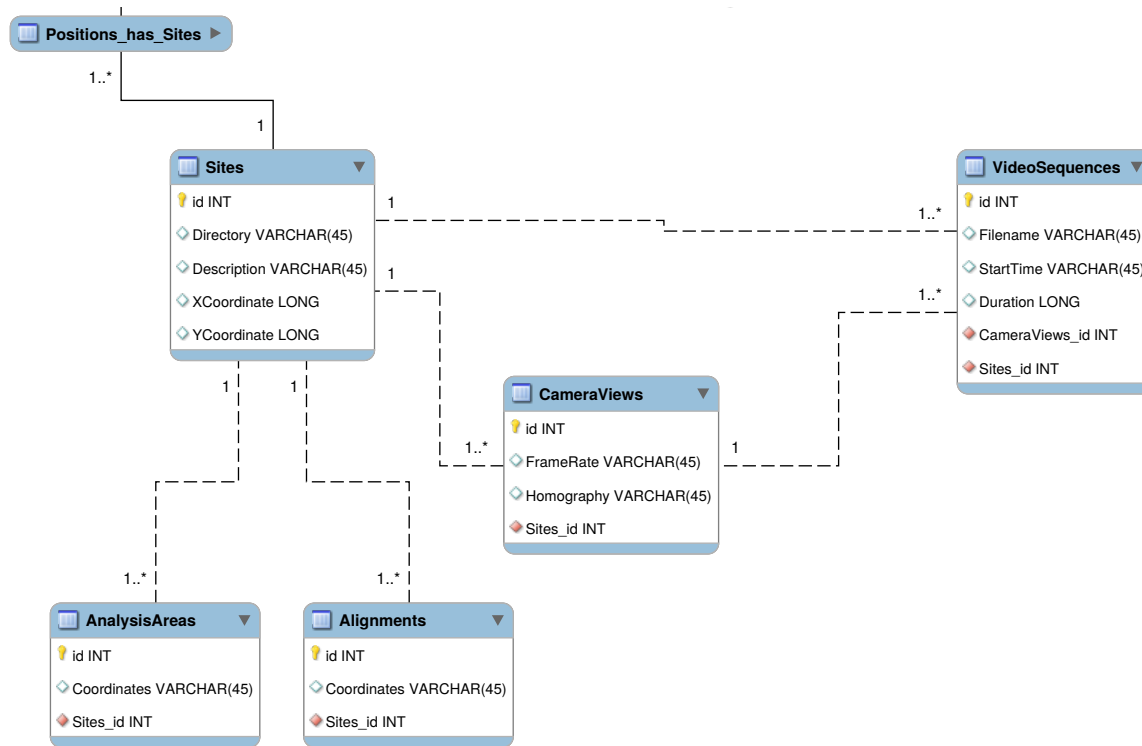


Figure 3.11 Metadata data model.

This metadata is stored in a single, portable, cross-platform, zero-configuration database file⁴ and the file is stored with the video data. The metadata implements six basic data abstraction layers, allowing for indexed access to specific portions of video and tracking data according to study needs:

- A list of **camera types**. This is a database of the different cameras used during data collection. It includes the following information: intrinsic parameters (for lens distortion correction), frame rate, resolution, and scaling parameters.
- A list of **video sequences**. This is simply a list of all video sequences within the database. It contains information about the video file: filename, start time, duration, global translations and rotations, and tracking configuration. The same video file can

⁴SQLite is used for the purpose of this implementation, as it meets these specifications.

have multiple entries with different tracking configurations and thus different sets of trajectories (advanced usage).

- A list of **camera views** which groups together a series of video sequences filmed together at the same time and place with the same camera type. The view is assumed to be static. As such, the computed homography and any optional mask are applied to all sequences within the camera view. A tracking configuration can also be applied to all sequences of the camera (default functionality).
- A list of **sites** with multiple corresponding camera views located at the same general geographical location (the site proper). Typically, this will be linked to a database of geometric parameters for site studies and site-based traffic analysis. Each site stores alignment data and also points towards aerial imagery and geographical coordinates of the origin point in world space—the lower-left corner of the aerial image.
- A list of **site analyses** with corresponding time range and analysis zones which act as temporal and spatial trajectory (or trajectory subsegment) selectors respectively. The implementation of the analysis zones selector is functionally identical to mask filtering, except that the tracked moving objects are permanently rejected with mask filtering, while analysis zone selection is used in the context of selecting specific trajectory for traffic analysis only. A site analysis also contains additional metadata, including time-sensitive factors, traffic expansion factors, virtual loops (which act as ordinary point-based traffic counters), and various data presentation options. For basic usage, each camera view will have one site analysis, but if the data permits, multiple site analyses can be performed on the same trajectory data from the same camera view.
- A list of **analyses** which cluster site analyses to produce aggregated results (for cluster comparison).

High-Level Interpretation (HLI) modules can extend the functionality of this metadata through the addition of extended tables or of abstraction layers in tandem (see section 3.4.3).

3.4.2 Alignments

Trajectory clustering is the first step in road user behaviour analysis. Trajectory clustering is the abstract representation of movements along prototypical paths through a scene, called **alignments**. This is the basis for relating the spatial position of trajectories to elements of road geometry and, in particular, the position of moving objects to traffic lanes, bike paths, and sidewalks (St-Aubin et al., 2015a).

Many approaches to trajectory clustering have been explored. While some methods are supervised (Schreck et al., 2008), many more are unsupervised, e.g. k-means (MacQueen, 1967) or hidden Markov models (Rodríguez-Serrano and Singh, 2012). Manual trajectory clustering is labour intensive and may be a source of bias, but it allows for tight control of scene description and analysis oversight. Unsupervised clustering is systematic but naive, as this form of clustering can only make use of trajectory data to infer spatial relationship. Manual clustering along a series of splines, called alignments, is chosen for its simple implementation and better retention of control over interpretation. A hybrid approach, which automatically refines the spatial positioning of manually defined alignments through traditional unsupervised clustering approaches, is considered for future improvements (Morris and Trivedi, 2008; Schreck et al., 2008).

The data structure used to represent alignments is a simple series of points with a beginning and an end, in the same direction of travel as the majority of movements along this path. A set of these alignments is the basis for the curvilinear position coordinate system, where the lane number represents the alignment’s index amongst the set.

3.4.3 High-Level Interpretation

tvaLib is designed to include High-Level Interpretation (HLI) modularity. While many traffic analysis tasks are common for all types of traffic studies, such as counting cars and obtaining lane-bound speed profiles, extracting and interpreting infrastructure-specific measures requires modular functionality to be included. For example, specific behaviour unique to roundabouts is expected at the approach of roundabouts. Evaluating this behaviour requires yielding post-encroachment time (yPET) measures for only those road users in the approach lane which are merging into the merging zone; this process requires additional metadata describing these points of infrastructure.

A basic modular framework is incorporated into tvaLib in the form of HLI to add study-specific metadata and analysis procedures to the general-purpose traffic analysis routines. A reference module is included by default with tvaLib and is designed to conduct analysis specific to roundabout geometry. This functionality is discussed in much greater detail in section 5.1.2 of the chapter on roundabouts (Chapter 5).

3.4.4 Technical Improvements for Large-Scale Data Processing

Data processing time varies depending on a wide variety of conditions including camera resolution, need for and degree of lens undistortion, presence of any masked portions of image




Figure 3.12 Example alignments and links for a typical roundabout. Trajectories projected to world space are in light grey.

space, and machine specifications. For a typical high-end consumer-grade machine available today, each hour of video typically requires one to two hours of compute time (single thread). Offloading calculations onto graphics processing units could significantly improve processing times, however this technology is not yet practice-ready. In any case, due to the large volume of data, additional features are implemented in the program architecture of tvaLib to more efficiently process data.

By default, tvaLib is designed to cache objects via memory serialization after initially loading objects from tracked feature files and applying filtering functions, because architectural and algorithmic constraints require that objects be loaded and unloaded from memory several times (in any case, this is useful during development and data experimentation). Experimentally, it is demonstrated that loading objects from a database is a relatively expensive

operation compared to memory serialisation. Loading from a cache reduces this time by about 95 % (i.e. about three minutes per sequence stored in database versus several seconds from cache). Note that objects are cached *before* targeting an analysis to a specific analysis zone—only mask filtering is applied to the cache. Also note that this caching is not intended for archiving or data transferral between machines; this is the reason trajectory data is stored to a portable, cross-platform, zero-configuration database file in the first place.

Many opportunities exist for massive parallelization of video processing to handle vast amounts of data. With the exception of machine-learned tasks (a.g. as will be covered in section 4.3.2), video sequences are stored individually and can be processed independently. Technically, motion pattern array (MPA) generation can be parallelized across analyses, however this is not recommended, as MPA generation is a memory-intensive task (upwards of 10GB of memory). Although parallelizing has been implemented in tvaLib for the calculation of collision-courses, parallelizing analysis tasks of entire sequences is the preferred operating procedure wherever possible, as it simplifies program architecture. Where possible, pre-calculated results to be compiled later using analysis scripts should be cached via serialization (this is handled automatically by tvaLib).

3.4.5 Tracking Performance

Tracking accuracy is evaluated by way of the MOTA (Bernardin and Stiefelhagen, 2008) using as ground truth manually annotated trajectories with the open source annotation tool from the Urban Tracker project (Jodoin et al., 2016). With this methodology, a perfect tracking score, where all positions of all road users are matched continuously within a given search distance ψ , yields a maximum MOTA of 100 %. However, MOTA can become negative if more false alarms are detected than there are ground truth objects.

As part of an effort to optimize tracking results, nearly 40,000 observations (instants) of 371 motor vehicles are annotated manually from training data collected at two prototypical sample sites: one site where objects are tracked at a distance of between 20 and 50 metres from the camera and one site where objects are tracked at a distance no greater than 20 m. An initial MOTA with $\psi \leq 3$ m is calculated for default tracking parameters; this offers modest performance (in the neighbourhood of 70 %). A tracking optimization is then performed using a genetic algorithm (similar to simulated annealing in Ettehadieh et al., 2015) to search for parameters that have maximized MOTA. After optimization, accuracy increased to 94 % (measured over the same ground truth used for optimization. These statistics are summarized in Table 3.1. Tracking parameter optimization converged in less than 24 hours.

Table 3.1 Tracking Performance

Site	Ground Truth Objects	Ground Truth Observations	Unoptimized MOTA (Default Parameters)	Optimized MOTA
#1-1-12 (far camera)	146	12887	0.685	0.944
#22-3-16 (near camera)	225	25011	0.744	0.853

While these results are very good for individual sites, they are arguably over-fitted. For true automation, it is desirable to find a single set of parameters that can be applied to any tracking conditions, or to an easily identifiable class of tracking conditions. To this end, a more thorough calibration attempt is then performed (in Morse et al., 2016) across five sites over a greater range of visibility conditions, camera, angles, and camera types. It is found that winter conditions have the largest impact on tracking optimization, requiring a substantially different set of tracking parameters. It should be noted that MOTA fails to capture the velocity smoothing half-width tracking parameter. Otherwise, most general tracking performance yields low 80 % MOTA values. These optimized tracking results are implemented on the final set of data used for model selection in section 4.6.2 and Chapters 5 and 6. The general summer parameters used for all data in the remaining chapters are presented in Table 3.2. None of the final data used for this project was collected during winter. Future work might examine effects of weather on tracking in more detail.

3.4.6 Quality Control and Error Correction

While Traffic Intelligence impacts trajectory quality primarily through a number of feature-based heuristics and also from velocity smoothing, tvaLib implements a secondary layer of filtering composed of several empirical filtering functions to detect and clean erroneous moving objects. While the aim of tracking optimization is to improve the overall accuracy of all tracked points (high-volume, low-impact quality control), the objective of this secondary layer of filtering is the preparation of data for adequate traffic analysis and the removal of specific problematic trajectories that are known to greatly skew traffic analysis results (low-volume, high-impact quality control).

As these functions modify moving objects directly, they must be run in a specific order. The functions should be applied in the order in which they are listed here:

- A function taht verifies the data integrity of moving objects loaded into memory and

Table 3.2 Optimized General Summer Feature Tracker Parameters

Parameter	Range	Description	Default Value	Calibrated Value
Feature Tracking				
feature-quality	[0-0.4]	Minimum quality of corners to track (unitless)	0.1	0.08
min-feature-distance-klt	[0-6]	Minimum distance between features (in pixels)	5	3.55
window-size	[3-10]	Distance within which to search for feature in next frame (in pixels)	7	6
min-tracking-error	[0.01-0.3]	Minimum error to reach to stop optical flow (unitless)	0.3	0.183
min-feature-time	[2-10]	Minimum time a feature must exist to be saved (in frames)	20	9
Feature Grouping				
mm-connection-distance	[1.5-3]	Distance to connect features into objects (in world distance unit (m))	3.75	2.69
mm-segmentation-distance	[1-3]	Segmentation distance (in world distance unit (m))	1.5	1.81
min-nfeatures-group	[2-4]	Minimum number of features per frame to generate a road user	3	3.17
Other parameters				
smoothing-halfwidth		Half-width of velocity smoothing window (simple moving average)	5	5

searches for data formatting corruptions. Data corruptions are rare, but are bound to occur in large volumes of data for a number of technical reasons related to software bugs, misconfigurations, terminated processes, or interrupted file transfers. This function simply checks that data is properly formatted following data structure (e.g. trajectory storage specification) and that data containers have sizes consistent with expected values.

- A function that applies a translation vector v to a point p (any basic point in world space as defined by equation 3.7), and obtains $T_v(p)$ followed by a rotation R of all moving object points in world space around the origin $(0,0)$ by an angle θ . This can be used to quickly realign trajectories to correct geometry, if the camera changes orientation slightly from one sequence to another. For significant changes in orientation

or repositioning of the camera, a full homography recalibration should be used instead.

$$T_v(p) = p + v \quad (3.15)$$

$$R(x, y) = (x\cos(\theta) - y\sin(\theta), x\sin(\theta) + y\cos(\theta)) \quad (3.16)$$

- A function that constrains a list of moving objects to a particular region of space by removing entire trajectories, sections of trajectories, or individual points of trajectories outside a given set of polygons. This process of removal is applied up to three times:
 - firstly, and optionally, regions in image space (e.g. without traffic) can, to improve performance, be designated to be ignored by the feature-based tracker;
 - secondly, a tracking mask is designated for the purpose of rejecting specific tracking regions yielding poor results (frequently found at the boundaries available to the feature-based tracker, as only partial feature tracking occurs at these); and
 - finally, a third time to be used as a moving object selector as part of the analysis zone.

Trajectories are assumed to enter and leave the polygons exactly once and are truncated successively from the start and end of the trajectory inwards. If a section of trajectory lies outside a polygon, but between two sections contained by the polygons, then that section will be retained as though it were effectively contained by the polygon(s) as well. This function has one empirical parameter which is intended to catch extreme cases where trajectory ends pop in and out of the polygon(s). If the ratio of removed points to existing points is greater than `bb_containment_threshold`, the trajectory is removed entirely. A further refinement determines entry and exit points of the polygon(s) within a 20-frame window to account for occasional noisy positions (not shown in algorithm 2). The function is presented in algorithm 2.

- A function that splits trajectories on key turn points (outlier vectors which deviate significantly from the trajectory via large and/or sudden change in speed defined by parameters `so_soft_maxSpeed` and `so_hard_maxSpeed`, and heading `so_max_angle`. These typically occur when features of different real objects are grouped in the same moving objects, or, in more extreme situations, when features from different real objects become mixed, especially when one moving object leaves the scene and another enters it nearby, immediately after. An example of a trajectory corrected in this way can be seen in Figure 3.13. The algorithm is presented in algorithm 3.

Algorithm 2 Filter to Bounding Box

```

1: Input: objects (a series of moving objects), polygons (a series of polygon objects each
   defined by a series of points representing its vertices), bb_containment_threshold.
2: Output: modified objects
3: function FILTERBOUNDINGBOX(objects, polygons, bb_containment_threshold)
4:   for object in objects do
5:     for point in object do
6:       if point not in polygons then
7:         Add point to list of rejected points
8:       if count of rejected points  $\geq$  bb_containment_threshold  $\times$  length of object then
9:         delete object
10:    for point in object do
11:      if point in rejected points then
12:        lower point = point break
13:    for point in reverse ordered object do
14:      if point in rejected points then
15:        upper point = point break
16:    truncate object to the range: [lower point, upper point]
   return objects

```

Algorithm 3 Split by Outlier Vector

```

1: Input: objects (a series of moving objects), so_soft_maxSpeed, so_hard_maxSpeed,
   so_max_angle.
2: Output: modified objects
3: function SPLITBYOUTLIERPOINTS(objects, so_soft_maxSpeed, so_hard_maxSpeed,
   so_max_angle)
4:   for object in objects do
5:     for velocity vector  $\vec{v}_i$  in object (less one) do
6:       if  $\|\vec{v}_i\| > \text{so\_hard\_maxSpeed}$  or ( $\|\vec{v}_i\| > \text{so\_soft\_maxSpeed}$  and  $|\text{angle}(\vec{v}_i) -$ 
    $\text{angle}(v_{i+1})| > \text{so\_max\_angle}$ ) then
7:         Split object into two objects at this instant
8:   return objects

```

- A function that removes moving objects with short trajectory lengths, i.e. whose trajectories are shorter than `te_config_min_traj_len`. These are almost always duplicate representations of the same real object (*oversegmentation*), noise, or by-products of previous filtering functions. For issues with partial false negative detection of real road users, these short trajectories can cause other issues with data interpretation techniques in Chapter 4 and 5.
- A function that detects oversegmentation explicitly. For each pair of coexisting moving objects, similarity and proximity over a contiguous section (e.g. 80 %) of the length of the joint trajectory is evaluated using minimum speed similarity $s_{f,min}$ (i.e. `td_minimumSeparationVelocity`) and minimum spatial proximity $d_{f,min}$ (i.e. `td_minimumSeparationDistance`). This is presented in algorithm 4. An alternative oversegmentation detection algorithm might implement an alternative motion similarity algorithm instead, e.g. LCSS or dynamic time warping (DTW).

Algorithm 4 Filter Tracking Duplicates

```

1: Input: objects (a series of moving objects), td_contiguity (base value of 0.8 or 80 %),
    $s_{f,min}$ ,  $d_{f,min}$ .
2: Output: modified objects
3: function DROPTRACKINGDUPLICATES(objects, td_contiguity,  $s_{f,min}$ ,  $d_{f,min}$ )
4:   for object in objects do
5:     for object2 in objects do
6:       commonTimeInterval = common time interval of object and object2
7:       if object is object2 then continue
8:       if commonTimeInterval is empty then continue
9:       if more than contiguity percentage of points at common times have a distance
   >  $d_{f,min}$  then continue
10:      if more than contiguity percentage of points at common times have a differ-
   ential absolute velocity >  $s_{f,min}$  then continue
11:      if object2 has more features than object then
12:        remove object
13:      else
14:        remove object2
   return objects

```

- A curvilinear transformation function that calculates and adds, for every moving object position, curvilinear position coordinate transformation as per algorithm 1 using

alignment metadata. With this information, it is possible to check for heading consistency. Without automated detection of road violations (e.g. driving against traffic), these tracking errors may be reviewed manually instead.




Figure 3.13 An example trajectory that jumps from one road user to another. The corrective action is to split the moving object at the outlying vectors.

The filtering functions are calibrated using several empirical factors which undergo basic calibration performed on a case-by-case basis via manual inspection of the tracking and interpretation results. Application-specific calibration of these parameters is still technically possible, and has been attempted, using the MOTA optimization tool. However, since the volume of corrections is low, and MOTA only measures spatial accuracy of trajectories and cannot account for *impact* on analysis, use of MOTA optimization on these filtering parameters is generally not recommended. The default calibration parameters used for the data throughout the remaining chapters are listed in table 3.3.

Table 3.3 Summary of Default Filtering Parameters

Parameter	Value	Units or range
bb_containment_threshold	0.9	[0,1]
so_hard_maxSpeed	6.0	m/frame
so_soft_maxSpeed	2.0	m/frame
so_max_angle	45.0	degrees
te_config_min_traj_len	20	frames
td_minimumSeparationDistance	8.0	m
td_minimumSeparationVelocity	0.08	m/frame
td_contiguity	0.8	[0,1]

3.4.7 Playback

Traffic Intelligence incorporates basic trajectory playback to display tracked trajectories over image-space. `tvaLib` improves this tool with a graphical user interface (GUI), an object and interaction explorer, world space projection coordinate display and visualisation tools, a timeline scrubber, and a time series display of various measures such as speed, curvilinear position, and interactions. A preview of this interface can be seen in Figure 3.14.



Figure 3.14 tvaLib GUI for trajectory, infraction, ground truth, and interaction playback.

This tool is vital in performing quality control of the feature-based tracker and filtering process, and helps the end-user visualize the analysis process and analysis outputs. For the case studies conducted in Chapters 5 and 6, random sites were chosen for data quality control and normal supervision of the analysis.

3.5 Summary

In this chapter, a complete, fully automated traffic data collection system is designed and presented. Such a system is designed to be integrated in a traffic engineering or road safety analysis study. From video data collected with a VCU, the system records continuous, microscopically-detailed road user trajectory data. Furthermore, it is designed to efficiently scale for large volumes of data collected at an unlimited number of sites, a technical innovation lacking in many proof-of-concept systems build in the past.

The core contribution of this work lies in the numerous technical innovations discussed throughout this chapter that permit the scaling necessary for a complete, automated, continuous trajectory data collection system, summarised in the analysis pipeline outlined in Figure 3.9. Additionally, general-purpose, semi-automated trajectory descriptors are implemented in the form of trajectory alignments and curvilinear coordinate transformation (to be used in subsequent chapters). Finally, basic validation of the trajectory data is performed through tracking performance optimization, quality control, and systematic error correction in the data. However, more sophisticated trajectory matching algorithms could benefit the system, particularly with respect to trajectory connecting (as in Sankararaman et al., 2013).

Although real-time or near-real-time analysis is not the stated goal of the system, near-real-time analysis may be feasibly implemented with this system via the proposed parallelisation of tasks. This is, in fact, even necessary in order to process in a timely manner the volume of data (big data) required for some traffic engineering projects.

An important limitation of this video-based traffic data collection system is its dependence on line of sight and on the quality of lighting conditions. However, no other external sensor technology exists that can reliably capture high-fidelity positions of road users continually. Practical options exist to limit these issues. For example, line of sight issues can be addressed with higher-angled VCU installations or even aerial drones (as in Puri, 2005; Coifman et al., 2006; Kanistras et al., 2013). Meanwhile, the video-based traffic data collection system is compatible with any type of video data, including optic sensors designed to operate in poor lighting conditions (e.g. thermographic optics) or even hyperspectral imaging solutions, i.e. data captured from multiple electromagnetic spectra simultaneously.

The microscopic road user data obtained is now ready to be interpreted using road safety-specific methods of analysis and geometry-specific measures of traffic flow as presented in Chapters 4 and 5 respectively. Other uses for this type of large volume microscopic data includes calibration of microsimulation and car-following models, automated traffic incident or infraction detection, and a variety of specialized traffic research applications.

CHAPTER 4 SURROGATE SAFETY FRAMEWORK FOR ROAD SAFETY ANALYSIS

In Chapter 2, section 2.2, a basic safety framework based on historical accident data was reviewed from literature. In this chapter, this framework will be expanded to include surrogate safety methods. However, given that surrogate safety methods, by definition, typically substitute historical accident data, there is a need for an alternative model that can predict the expected number of collisions from the available data.

In the previous chapter, large quantities of high-resolution traffic data were collected. The purpose of collecting such large amounts of data at such a fine level of detail will become apparent in this chapter; these are required for any type of systematic surrogate safety methodology. At this stage of the work, trajectory validation and error correction has been completed, and all resulting moving objects are now assumed to be road users, described by the trajectory of their centroid and other moving object characteristics.

Here, a complete framework for objectively-defined surrogate safety method measurement and analysis is presented. Recalling the discussion of section 2.3, surrogate safety methods are a proactive family of approaches to quantifying various aspects of road safety using the indirect safety indicators, also known as surrogate safety measure (SSM)s in this context specifically. In this framework safety indicators defined earlier in the literature (e.g. speed, yPET, and TTC) are revisited and newer SSMs are defined (pPET) or adapted from older surrogate safety methods (e.g. MPA for TTC calculation). The majority of this chapter is methodological in nature, though some exploratory analysis is performed towards the end of the chapter to test, compare, or demonstrate the performance of safety indicators. The final objective of this chapter is to carry out a comparative analysis of a set of independent safety indicators for further use in studies, and to be used directly in the case study of roundabouts presented in Chapters 5 and 6.

4.1 Introduction to Surrogate Safety

A **surrogate safety measure (SSM)** is a type of traffic data intended to be used to describe and quantify the various aspects of road safety indirectly, that is, a method which uses some source of data other than historical accident data (e.g. accident reports recorded and compiled by transportation authorities). Thus, SSMs are a specialized type of safety indicator. As discussed earlier in Chapter 1, the appeal of surrogate safety methods is their

proactive nature when compared to historical accident analysis methods: surrogate safety methods do not require 10-year-long studies and they do not subject the public to unknown levels of risk with untested traffic designs for nearly as long as a historical accident analysis requires. SSMs come in many different forms, and the seasoned reader may even be aware of some examples of SSM, although possibly under a different name, for example: speed.

On the practical side, with the exception of applications of surrogate safety methods in studies using driving simulators and naturalistic driving studies, all SSMs are designed to be externally observable, that is to say, they are visible to a third party observing road users from the side of any road.

The general principle behind surrogate safety methods is that precursors or preconditions of relevant aspects of road safety exist for distinct traffic events. These precursors or preconditions are measured using corresponding SSMs, each of which models an aspect of road safety, i.e. collision probability or collision severity or both. Thus, for the same discrete traffic event, given a precondition that is a “worse” outcome for road safety, the expected road safety result of that traffic event is an increase in probability of collision ($P(\textit{collision})$) or a more severe collision outcome (collision severity). In this way, it is established that SSMs should be defined such that they are proportional (or correlated):

$$P(\textit{collision}) \propto f_1(SSM, \dots) \quad (4.1)$$

or

$$\textit{collision severity} \propto f_2(SSM, \dots) \quad (4.2)$$

where *collision severity* measures any collision outcome (conditional to a true collision occurrence) in units of PDO, injury, or fatality, or some compound of the three, using a common scale (such as cost, see National Safety Council, 2013). The relationship between SSM and $P(\textit{collision})$ is difficult to quantify. There are many types and causes of collisions, and the mechanics involved are more complicated and less well understood. The relationship between SSMs and $P(\textit{collision})$ has been quantified in the literature to varying degrees, though no evidence that SSMs and $P(\textit{collision})$ are not related has been found. One early study (Hydén, 1987) produced conversion factors calibrated to convert the number of serious traffic conflicts directly into an expected number of collisions during the same period. Nevertheless, despite the small availability of quantitative evidence, there is overall qualitative consensus

that SSMs are related to $P(\textit{collision})$ in some way. As discussed earlier, the relationship between speed as a SSM and collision severity is well documented. This chapter presents a framework for the development of formally-defined, objectively-measured SSMs capable of modelling collision probability mechanics specifically, though calibration of these specific models remains for the future, as the historical accident data collection required to calibrate these models will take many years to collect. In the meantime, qualitative interpretation of SSMs is predominant in surrogate safety methods usage.

Furthermore, given more preconditions or discrete traffic events of the same nature—i.e. exposure—the total number of collisions at a given site is expected to increase. Thus, this surrogate safety method is expanded to provide a model of *expected number of collisions p.a.* at a given site for a given period of time in place of historical accident analysis. This model is the product of the probability $P(\textit{collision})$ of any predefined discrete traffic event resulting in a collision with annual exposure to those traffic events (as in Hauer, 1982):

$$\textit{expected collisions p.a.} = P(\textit{collision}) \times \textit{discrete traffic events p.a.} \quad (4.3)$$

By construction, those characteristics or situations that make a good collision-probability-based SSM are necessary precondition to resulting collisions, although only a subset of these situations with observable SSMs bear collisions. In other words, a SSM is not a guarantee of collisions, but it should be well chosen such that that SSM is observable at that collision. How those situations are defined depends largely on the framework adopted by the surrogate safety methods and the specific SSM being evaluated.

Finally, by modelling *expected number of collisions p.a.* together with collision severity (either with historical accident data or *collision severity* from SSMs) as in equation 2.1 an *expected PDO p.a.* (or expected fatality rate, etc.) can be obtained in order to make road safety management decisions, e.g. hotspot analysis.

4.1.1 User Pairs and Interaction Instants

As discussed in Chapter 2, gross aggregated measures of traffic exposure, such as total traffic flows, serve as crude preconditions to collisions, e.g. as controlled for in historical accident analysis. However, much better preconditions to collisions (e.g. discrete traffic events with collision potential) can be devised. For a collision to ensue, it is not enough that many road users arrive at a traffic scene; they must arrive simultaneously and generate interactions

(Saunier et al., 2010) between road users.¹ Thus it follows that measures of exposure can be improved using more detailed, microscopic traffic data which incorporate time-dependent observations. The most accurate time-dependent data for this task comes from computer-vision-based traffic data collection systems, as all road user trajectories are tracked near-continuously (several times per second).

As was shown in Figure 3.8, trajectories extracted from video data are sufficiently continuous to generate a new measure of exposure: the **user pair**, defined conceptually as a set of two (and exactly two) road users, or theoretically one road user and one stationary object or obstacle within the road, which coexist in time and with enough spatial proximity that some meaningful collision probability may be generated, i.e. within the area of interest for study. While traffic accidents involving more than two road users are relatively common, these are usually chain-reaction events that fall under the scope of modelling collision severity. Furthermore, it is also possible for a road user to be involved in multiple user pairs simultaneously. Simultaneous interaction between more than two road users is intrinsic to the model.

This is not a unique approach to improving exposure measures. Previous work in the area of cyclist safety has used intersection flow data disaggregated to the type of traffic movement (right-turn, left-turn, conflicting with cyclist movements etc. in Miranda-Moreno et al., 2011a; Strauss et al., 2013) (lateral conflicts at signalized intersections Midenet et al., 2011). The proposed framework aims for a more generalized approach to disaggregated exposure measures.

By definition of exposure, as the number of user pairs increases, the total expected number of collisions is expected to increase, but at a decreasing rate, as per the “safety in numbers” effect. Thus, this measure of exposure already serves as a very basic, though admittedly not particularly useful, SSM. Note that Hauer explicitly suggested that exposure is diametrically opposed to traffic conflict, suggesting that the $P(\text{collision})$ of traffic conflicts should be chosen *after* accounting for exposure (Hauer, 1982). To be clear, in this framework, user pairs serve to estimate *discrete traffic events p.a.* as per equation 4.3. User pairs are part of the classic surrogate safety method hierarchy (e.g. the safety pyramid suggested by Hydén, 1987) of discrete traffic events. The $P(\text{collision})$ will be further investigated using more sophisticated SSMs and directly observed SSMs.

More precisely, the user pair is stored as a pair of road users occupying a given shared space over a common time interval. This shared space is defined as the **analysis zone**, chosen so as to be shared by all road users to be studied, large enough to include road user approach

¹Arguably, this is not the case for *some* types of collisions involving a single road user.

vectors, i.e. sufficient observations to determine speed and heading upon entering the analysis zone, of road users, but small enough that all road users are at all times within “interaction-distance,” i.e. within a distance from each other that trajectory negotiation is necessary to avoid collision. Some example applications include a small section of straight road (e.g. 100 m in length) with multiple lanes and dense traffic (St-Aubin et al., 2013b), or an ordinary traffic-light-controlled intersection, or a highway junction, among other examples.

The coexistence of user pairs in time involves a common time interval where both user pairs are moving jointly and in proximity (bounded by the analysis zone) to each other. This forms a time series of interaction instants,² i.e. a series of instants of spatial proximity between the two user pairs defined by measures such as centroid-to-centroid as well as volumetric proximity, and relative velocity.

In any case, user pairs serve as the base traffic event for most of the collision-probability-based SSMs to follow. Thus, in this application of equation 4.3, $P(\text{collision})$ is defined as the probability that a user pair (the discrete traffic event) results in a collision between both road users, and *discrete traffic events p.a.* can be observed or estimated from an observed number of user pairs in a year. Experimentally, it can be shown that the number of road users increases quadratically and almost exclusively with traffic volume, and this, regardless of the number of lanes present in the analysis zone, as demonstrated in Figure 4.1. Thus, if yearly traffic volumes are known, predicting number of user pairs in a year is a relatively trivial exercise.

This relationship between hourly user pairs U_{pairs} and hourly traffic volumes Q is modelled from the datasets used for the comparison of Québec and Swedish roundabout merging zones presented in Chapters 5 and 6 for all types of lane configuration using a simple quadratic expression:

$$U_{pairs} = 0.0014Q^2 + 0.0311Q - 4.5397 \quad (4.4)$$

The quadratic nature of this relationship can be simply explained by the fact that a larger hourly traffic volume increases the probability that more than one road user is present at a time in the analysis zone and that for every additional road user present in the analysis zone, an additional number of user pairs is generated, equal to the average number of existing road users.

²Terminological note: a conscious choice is made to avoid the term *interaction* (alone) to refer to time series of interaction instants, as the meaning is vague in English and ambiguous across fields of science, and even in different traffic engineering contexts. The term user pairs is more explicit and preferred.




Figure 4.1 Hourly user pairs versus flow rate measured at different times of the day at twenty different roundabout merge zones split between Québec and Sweden.

4.2 Directly Observable Surrogate Safety Measures

Directly observable SSMs are the most basic types of SSM and feature immediately observable behaviour. These are measures recorded, aggregated, and presented directly from instrumented data, e.g. trajectory data. Directly observable SSMs make no assumptions about possible road user collision outcomes—virtually no road users enter collision within the data set, as mean times between collisions are drastically higher than data collection durations. These SSMs tend not to model collision mechanisms explicitly.

Prevalent examples of directly observable SSMs in the literature include

- **Speed** is measured directly from road user trajectories or any other suitable speed recording instrument, and its effects on collision severity is very well understood in the literature (Fildes and Lee, 1993; Elvik et al., 2004). Use of recorded speed as a SSM is wide-spread, especially for modelling collision severity, however its exact significance with respect to collision probability is debatable, or at the very least contextual. Some researchers have debated the usefulness of speed measures for modelling collision probability (Hauer, 2009), though it is clearly demonstrated that speed affects operational characteristics such as braking time, and can account for the increased number of collisions observed in adverse weather conditions, where braking time is significantly longer for similar high speeds (Pisano et al., 2008).

Speed is the norm of the velocity vectors which can be obtained (derived) continuously from positions (see section 3.3.2) along road users' trajectories. For purposes of analysis, these observations are aggregated in a number of different ways depending on context, and include

- mean or median speed of each trajectory;
 - maximum or minimum speed of each trajectory; or
 - speed measured at a predetermined location (following context).
- **post-encroachment time (PET)** (Allen et al., 1978) represents a time lapse between the arrival of two successive road users on overlapping trajectories (Gettman and Head, 2003). It is measured directly from each road user's true movement. This measure holds qualitative value as measuring “near-miss” traffic events (as in Kraay et al., 2013). However, its effect on collision probability is less well studied quantitatively than some other measures. Conceptually, *PET* is measured as follows:

$$PET = |t_{p_2=p_{1\cap 2}} - t_{p_1=p_{1\cap 2}}| \quad (4.5)$$

where $p_{1\cap 2}$ is the point of intersection of the trajectory of each road user in a user pair and $t_{p_i=p_{1\cap 2}}$ is the arrival time of the road user $i \in 1, 2$ at the point of intersection $p_{1\cap 2}$.

Although the goal of the surrogate safety method of this chapter is to be as general a framework as possible, some aspects of road user behaviour are unique to specific road infrastructure designs. This is illustrated in the case of roundabouts where road users approaching a merging zone are expected to yield to other road users. To account for this, a special type of PET, redefined as the **yielding post-encroachment time (yPET)**, is proposed, which measures the lead and lag gap acceptance of the road user upon merging (effectively, a special type of lane change). Specific details about its implementation and calculation are provided later, in section 5.1.2.

4.3 Collision-Course-Based Measures

Collision courses are based on the idea behind the original traffic conflict definition: a collision course is defined precisely as any interaction instant between two or more road users that has some non-trivial collision probability in the near future.

Collision-course-based SSMs are distinguished from direct SSM by, instead of measuring a known, observable state (e.g. speed) with an uncertain outcome (the probability or severity of a collision), they measure an uncertain state (a collision course), with an uncertain outcome.

The true outcome of every interaction instant (collision or no collision in reality) is always known. Since ordinary road user behaviour data is used and collisions are rare events, nearly all observed trajectories do not yield real collisions. Instead, collision courses are formulated using predicted trajectories. Predicted trajectories are trajectories generated from a specific set of expected motion models known as motion prediction following a specific set of circumstances (*initial conditions*) which reflect both typical road user behaviour and specific collision precursors as well as expected road user behaviour. They model could-have-been scenarios, i.e. similar scenarios induced by the prevailing traffic and geometry conditions, but where the outcome is a collision, because a few unmeasured factors (e.g. distraction, visual obstruction, etc.) result in failure to avoid a collision course. Various hypotheses about trajectory motion prediction exist (St-Aubin et al., 2015b) and will be tackled in greater detail in the next two sections.

An example is provided to illustrate these concepts. Suppose a driver is found to be on a collision course with another road user, a cyclist, who does not see the driver for some reason or another. For the time being, the collision course will be considered obvious and imminent, as motion prediction models will be covered later. Drivers are/ expected to correct their collision course to avoid colliding with the unsuspecting cyclist. If the driver is distracted and fails to detect the collision course at all, it can be easily assumed that the result will be a collision. Furthermore, even if the driver does detect the collision course but there is insufficient time to break or evade it, then it can easily be assumed that the collision course will still result in a collision. Ideally, modelling this road user behaviour would enable the determination of the outcomes of these collision courses.

Unfortunately, as outside observers, this in-vehicle data is very difficult to obtain for a variety of ethical and logistical reasons. Instead, surrogate safety methods should substitute this information with safety indicators, i.e. some substitute measure with the predictive power of a road user's chances of correcting a collision course. This hypothesis is made on the basis that these events (distraction, occlusion, mechanical failure, etc.) are probabilistic and predictable with a statistical distribution. As will be shown later, this assumption somewhat limits conclusiveness of results, but this challenge can be overcome as the sophistication of the model increases.

Collision-course-based SSMs are always predicted instantaneously, if they exist, at each interaction instant, from the predicted trajectories of the two road users of any pair. Predicted

trajectories that intersect form the crossing zone. From the crossing zones where both road users are expected to arrive simultaneously, a collision point (CP) is created instead. In concept, CPs are a subset of crossing zones, however, these are handled exclusively in practice.

Conceptually speaking, each CP and crossing zone has, in addition to information about the user pairs involved, (i) a location, p , where the predicted trajectories intersect in *both* time and space; (ii) a probability of occurring $P(CP)$, defined by the probability of the intersection of each predicted trajectory, i.e. the product of the probability of each predicted trajectory arriving at point p ; and (iii) a SSM, i.e. a safety indicator. Thus, the collision point predicted at interaction instant t_0 has x, y coordinate components and a predicted probability of collision $P(CP)$ at time Δt after interaction instant t_0 . This is illustrated in Figure 4.5.

At any one interaction instant, there may be any number of CPs and crossing zones, or none at all. Some surrogate safety methods aggregate all CPs and crossing zones into a single safety indicator at a specific interaction instant by assuming that collision events are mutually exclusive and by taking the union of predicted collision events $P(\cup CP)$, which in this case evaluates to

$$P(\cup CP) = \sum P(CP) \quad (4.6)$$

and then by generating a weighted average indicator among the CPs and crossing zones (see the example in section 4.3.2). Spatial information is lost during this operation.

The CP's probability value is the collision probability for that specific interaction instant, assuming that the predicted trajectory models are valid, i.e. that the road users continue on their collision course. In practice, *road users react to collision courses dynamically and adjust their trajectories to actively avoid collision courses*. What this means is that there is an unaccounted probability that the predicted trajectory models are invalid and that evasive action occurs and succeeds ($P(evasion)$) in response to a collision course. $P(evasion)$ may depend on many factors such as road user attention, visual obstruction, etc. As stated before, these are not observed, but they, or their effects on $P(evasion)$, may be substituted by the safety indicator that is attached to the CP, and any other relevant $factor_k$ possible such that

$$P(evasion) = f(safety\ indicator, factor_k, \dots) \quad (4.7)$$

Thus, the true probability of collision $P(collision)$ is a product of the total calculated probability of collision ($\sum P(CP)$) and the probability of failing to evade the collision course.

$$P(\text{collision}) = \sum_i P(CP_i) \times (1 - P(\text{evasion}_{CP_i})) \quad (4.8)$$

So far, few, if any $P(\text{evasion})$ models exist. This is essentially an error term, with parameters (safety indicator) to fit some collision conversion factor model. Thus, as the sophistication of $P(CP)$ increases, the uncertainty handled by some error model $P(\text{evasion})$ is expected to decrease. This has been proposed in a very limited capacity in (Mohamed and Saunier, 2013).

SSMs of interest for the new surrogate safety method are

- **Time-to-collision (TTC)** is a very popular SSM which measures the remaining time before a collision occurs between two road users on a collision course. It is based on the principle that, as TTC decreases, the chance that at least one road user reacts to or engages in evasive manoeuvres before a collision ensues decreases dramatically (Hayward, 1971), since human reaction time lies within a relatively narrow range. However, reaction time is arguably dependent on sight distances generated by the collision course itself. Under “ordinary” conditions, average reaction time is commonly measured at 1.5 s in the literature (Hydén, 1987; Green, 2000) but will vary from individual to individual, and under different circumstances. All CPs necessarily have a TTC measure.

TTC can be measured by a number of different methods. The basic application of TTC, presuming motion following constant and linear velocity, takes the form

$$TTC = \frac{\|\vec{d}\|}{\|\Delta\vec{v}\|} \quad (4.9)$$

where \vec{d} is the distance vector separating two road users and $\Delta\vec{v}$ is the differential velocity vector between the two road users.

- **PET** is the time difference between the successive arrival of road users on observed, interacting trajectories at a common point (Allen et al., 1978).
- **predicted post-encroachment time (pPET)** is measured similarly to PET (the time difference between the successive arrival of road users on observed, interacting trajectories at a common point), however it is measured continuously, for each interaction instant, using predicted instead of existing predicted trajectories. All crossing

zones have a pPET measure (Mohamed, 2015). pPET is sometimes also referred to as **gap time (GT)** (American Association of State Highway and Transportation Officials, 2010).

Except for PET, which will be covered in the form of yPET as a yielding-specific SSM, in Chapter 5, the remainder of this work will focus on TTC, as it is one of the most widely-used and most generalizable SSMs. The hypothetical evolution of TTC on a collision course time series is illustrated in Figure 4.2. This evolution was first proposed by Hayward (Hayward, 1971) and subsequently reworked by Hyden and then Laureshyn (Hydén, 1987; Laureshyn, 2010), though this version has now been updated to include theoretical road user behaviour processes to model evasive action mechanisms.

During a collision course with TTC, road users must, in order of sequence, (i) perceive the collision course, (ii) react to the collision course by making a decision to initiate an evasive action, and finally (iii) execute the evasive action. When road users do not react dynamically to a predicted collision course, a collision ensues at $t = t_0 + TTC$ (by definition of TTC and collision course).

Note that the expected evolution curve decreases at a rate of one second per second for every instant where road user control input, e.g. driver action, is not initiated. See Figure 4.2 for an illustration of this process. This normally continues until the perception and decision phases end and evasive action is initiated on the part of the road user. Alternatively, this change in slope of TTC can happen if the collision course is altered for some other reason, i.e. a factor that the motion prediction does not take into account. At this stage, the collision course does not end immediately. An additional time-to-collision adjustment ($TTC_{A_{EAT}}$) is necessary, accounting for evasive action time (EAT), before the most critical possibility of collision for the user pair is averted, i.e. before the lowest TTC value, TTC_{min} , is reached. This TTC_{min} is effectively the margin of safety between all potential collision courses of the user pair.

Habitually, in the literature, if TTC_{min} falls below typical road user reaction times (e.g. 1.5 s (Hydén, 1987; Green, 2000)), the collision course is considered serious, or “doomed,” as the time remaining before a collision and that is required for evasive action falls below normal reaction times for road users. However, it can be argued that this state of being “doomed” can occur much sooner (e.g. at a time greater than 1.5 s) if EAT is significant. Recall that EAT is additional time necessary to avoid the critical portion of a collision course, i.e. TTC_{min} , after a road user begins to act. EAT sometimes depends on the collision course mechanisms, operating speeds, and possibly other factors. For example, in the absence of any room to steer out of the way of a collision course evasive action may rely on braking time

exclusively, and this may be problematic if speeds are high. Modelling EAT and resulting TTC_{EAT} is complex and is left as future work. Some early attempts at studying evasive action have been performed (in Mohamed and Saunier, 2013).




Figure 4.2 Hypothetical evolution of TTC over time for an ordinary collision course with road user evasive action (no collision).

What is interesting about TTC is that $P(evasion)$ may be hypothesized to increase as TTC increases (and possibly alongside other factors), probably non-linearly, especially because human reaction time is variable (within a small range). How this occurs is up for debate. Some researchers effectively use a $P(evasion) = 1$ if $TTC_{min} > 1.5$, and $P(evasion) = 0$ otherwise, as in SEC methodology. Meanwhile, others prefer to use a safety continuum comparison (SCC) (for example, Zheng et al., 2014b), e.g. where $P(evasion)$ increases rapidly up to 1.5 s and then approaches 1 as TTC approaches infinity. A hypothetical negative

exponential model with some calibration parameter α is proposed in (St-Aubin et al., 2012):

$$P(\text{evasion}) = 1 - \frac{1}{e^{TTC\alpha}} \quad (4.10)$$

SEC and SCC will be discussed in greater detail in section 4.6 regarding SSM evaluation.

4.3.1 Context-Free Motion Prediction

The classic approach to motion prediction is constant velocity, a context-free motion prediction. This class of motion prediction makes assumptions about road user motion, sometimes also referred to as kinematic motion prediction (Mohamed, 2015). This often involves the use of static hypotheses about road user behaviour and ability, encoded in prediction parameters. constant velocity motion prediction typically does not model specific contextual factors in road geometry and traffic conditions. These methods typically project trajectories using velocity information available at the instant of motion prediction.

- **Constant velocity:** This is the simplest and most classic hypothesis of motion prediction. It assumes that objects continue moving without acceleration in straight predicted trajectories. It is a narrow interpretation of the classic definition of a conflict “with movement remaining unchanged” (Amundsen and Hydén, 1977). Constant velocity simulates a situation in which control input ceases and road users are more or less subject to Newton’s First Law (the net forces of a vehicle are assumed to be zero or negligible). Constant velocity generates at most one collision point or one crossing zone, if any (in the case of diverging velocity vectors).

Constant velocity has no prediction parameters. If the interaction instant is not divergent, a crossing zone exists at the intersection of the straight lines formed by each of the two road users’ velocity vectors, \vec{v}_1 and \vec{v}_2 , passing through its centroid. This crossing zone is at a linear displacement of \vec{d}_1 and \vec{d}_2 away from each road user (themselves separated by a distance \vec{d}).

Furthermore, if this crossing zone is a collision point, a collision course is formed with a single corresponding TTC_{cvt} value. A crossing zone becomes a collision point when both road users arrive at the crossing zone simultaneously or near-simultaneously (accounting for road user volume), that is, $\frac{\|\vec{d}_1\|}{\|\vec{v}_1\|} = \frac{\|\vec{d}_2\|}{\|\vec{v}_2\|}$. Thusly, equation 4.9 is expanded as follows:

$$TTC_{cvt} = \frac{\|\vec{d}\|}{\|\vec{v}_1\|} = \frac{\|\vec{d}_2\|}{\|\vec{v}_2\|} = \frac{\|\vec{d}\|}{\|\Delta\vec{v}\|} \quad (4.11)$$

- **Normal adaptation:** A motion prediction method which modifies the hypothesis of collision course so that only non-evasive action control input is received, thus injecting small amounts of random, normal variation in predicted trajectory acceleration and orientation (Mohamed and Saunier, 2013). This input can be assumed to be either conscious or unconscious. This control input is modelled (as drawn from Mohamed and Saunier, 2013) using triangular distributions of acceleration and heading change which take a set of prediction parameters. Thus, the velocity vectors \vec{v}_1 and \vec{v}_2 are modified at each prediction step i with triangular distribution prediction parameters for acceleration $a = \pm 4.2 \text{ m/s}^2$ and heading $h = \pm 0.2 \text{ rad/s}$.
- **Point Set Prediction:** This is a special application of other motion prediction methods, where each position of the characteristic tracking features of road users is used instead of the aggregated position of each road user, in order to have robustness against micro-corrections caused by a continually changing number of tracking features (Mohamed and Saunier, 2013) and to work with road user volume. This method is also known as **Set of Initial Positions** (Mohamed, 2015).
- **Evasive Action:** This method of motion prediction is a stress test of a driver's ability to evade a collision, using extreme distributions of acceleration and steering (Mohamed and Saunier, 2013) in order to estimate $P(\text{evasion})$ (the inverse of which is called probability of unsuccessful evasive action in Mohamed and Saunier, 2013).

In this chapter, constant velocity and normal adaptation motion prediction will be reused as a baseline for evaluating the proposed methodological improvements to motion prediction methods, as constant velocity is the most popular implementation in the literature and normal adaptation is readily available.

4.3.2 Prediction Based on Motion Pattern

Motion patterns are empirical models of possible object movement (Morris and Trivedi, 2008). Applied to motion prediction, they are an approach whereby, through training data, methods are learned that predict the expected motion of a road user under normal circumstances toward any destination in (x, y) space Δt s into the future. Motion-pattern-based motion prediction methods have been implemented extensively outside of SSM applications. For example, the expectation maximization algorithm and HMMs are used to classify current, and predict future, pedestrian movement in (Bennewitz et al., 2005a).

These methods may be deterministic or probabilistic, or some combination of the two. Patterns are automatically learnt from observed road user trajectories through observed be-

haviour data of known trajectories of other road users experiencing similar conditions, such as speed and origin. Furthermore, unlike context-free motion prediction, motion-pattern-based motion prediction models can account for a wide variety of geometry-related elements and traffic conditions: the initial conditions.

As with some of the more sophisticated context-free motion prediction methods, multiple overlapping predicted trajectories are virtually guaranteed to exist, because multiple destinations are involved for multiple road users. The distinction between overlapping predicted trajectories and collision courses may be determined probabilistically. Thus, the advantages of motion pattern motion prediction lies in significantly increasing the modelling fidelity of naturalistic road user behaviour, which can be crucial for road safety and road user behaviour analysis applications.

One challenge presented by all motion pattern motion prediction methods is the uneven distribution of sample sizes between initial conditions present in most motion pattern training data. Motion prediction of outlying behaviour tends to be based on fewer samples than ordinary behaviour, due to outlier data. This can be demonstrated in the distribution of sample sizes for different initial conditions, shown in Figure 4.3. In this example, the distribution of sample sizes for each set of initial conditions follows the Pareto principle, i.e. 20 % of initial conditions are trained from 80 % of the trajectories, while the other 80 % of initial conditions are trained from only 20 % of the trajectories. However, this also means that motion prediction for the majority of trajectories is based on the largest samples, assuming enough data is collected overall. Motion pattern motion prediction is thus sensitive to outlying behaviour, but this issue can be mitigated by collecting large quantities of data.

Another problem with motion pattern motion prediction is the limited time horizon available for motion prediction. A time horizon is the maximum time that a trajectory can be projected into the future. With context-free motion prediction, this time is capped arbitrarily at about 5 to 10 s, as the hypotheses used to make context-free motion predictions longer than 5 or 10 s into the future hold little value—first, collision courses predicted that far into the future are unreliable, and secondly, it is unreasonable to expect that road users are incapable of executing an evasive manoeuvre in a span of time that long. However, with motion pattern motion prediction, an additional constraint exists on the span of prediction: the dwell time of road users used to calibrate the prediction model within the analysis zone. Since motion pattern motion prediction methods are typically trained only from road users within the analysis zone, only trajectories originating and destinating within the analysis zone can be modelled (effectively a shorter time horizon based on dwell-time within the analysis zone). This causes boundary issues along the analysis zone.

The effective time horizon used in the analysis of analysis zones presented in the remaining chapters is effectively 5 s, which is acceptable in comparison with context-free motion prediction issues of reliability. During statistical analysis, context-free motion prediction's time horizon is capped at 5 s in any case. To address this time horizon problem in the future, the analysis zone should properly encapsulate the area to be studied. Larger data sets that extend outside the analysis zones could provide larger time horizons. With current data sets, the time horizons of motion pattern motion prediction can be pushed a little bit further by including trajectories not only from the analysis zones, but also from the mask area. With trajectory connecting, this problem can be solved almost completely, assuming multiple cameras are used and trajectory coverage is good in relation to the analysis zone.

Motion Patterns Using Prototype Trajectories

A form of motion pattern motion prediction has been developed, which is the unsupervised prototypical road user movement clustering of trajectories by trajectory similarity.

These methods have been implemented in practice for SSM measurement, primarily based on LCSS (Saunier et al., 2007). LCSS particularly is recommended for its noise insensitivity (Vlachos et al., 2005). A practical application of prototype-based motion pattern motion prediction has been demonstrated by Saunier *et al.* (in Saunier et al., 2010) and by Mohamed (in Mohamed, 2015), using trajectories extracted from computer vision.

Discretized Motion Pattern Motion Prediction

Discretized motion pattern is a new motion prediction method proposed as a way of simplifying and generalizing prediction contexts and improving prediction performance. In discretized motion pattern motion prediction, motion probabilities are trained discretely (for ranges of initial conditions) and the data is stored in a multi-dimensional array, called a motion pattern array (MPA) (St-Aubin et al., 2014). In addition to time and the three dimensions of space used to encode destinations, the array has an additional dimension for each initial condition (motion context) parametrized, which dimension is used to identify corresponding similarity for road users whose trajectories are being predicted. Each additional dimension of initial condition is discretized in order to map it to a spatial distribution of the probability of being in a given discretized location in space at a future time. The current approach used in this research models a probability of destinating at position x, y , given speed $\|\vec{v}\|$, curvilinear distance S , and lane number $lane$ at prediction time t_i , though the number of parameters that may be added is limitless, provided enough training data is available to generate significant populations for each set of initial conditions.

$$MPA : (\|\vec{v}\|, S, lane, t_i) \mapsto P(x, y) \quad (4.12)$$

For example, road users with an initial speed $\|\vec{v}\| = 52 \text{ km/h}$, located $S = 5.4 \text{ m}$ downstream of the start of $lane = 1$, will have their future position on an $m \times n$ grid of x, y positions predicted probabilistically, to the nearest 1 m by 1 m square, at time $t_0 + t_i \times \Delta t$ for $i = 1..h$ in the future, where h is the maximum number of time steps, forming the time horizon t_h . This prediction is made from only those other road users with similar initial conditions: travelling at a speed of between 50 km/h and 55 km/h, and only if they are located between 5 and 7 m downstream of the start of $lane 1$. The sizes of m and n are chosen so as to cover the entirety of the analysis zone using the given prediction grid cell size, e.g. a 1 m by 1 m square. Thus, in this example, there is a total of six MPA dimensions.

In addition to the other advantages common to all motion-pattern-based motion prediction methods, discretized motion pattern motion prediction has the advantages of being simple and flexible to implement (in the sense that it can easily parametrize any externally-observable motion context). However there are technical challenges in its implementation and in data sampling, and the resulting predicted motion is discrete in both space and time (in this case, accurate to the nearest 1 m by 1 m square, which is slightly smaller than a typical motor vehicle or bicycle).

One practical problem specific to MPAs is that they become very large very quickly, especially as the range of initial conditions increases. In the above example, the MPA generated uses a 60×60 grid³ to encode destinations, an h of 100 increments, 5 originating lanes, and a range of 10 bins to categorize each of the remaining initial conditions (speed and curvilinear distance along each lane). Thus a $60 \times 60 \times 100 \times 5 \times 10 \times 10$ array is created, containing 180 million cells.

The solution to this problem is to use sparse array data structures to represent the data, as the vast majority of cells in the MPA are empty. This is because road users either do not have physical access to a large number of the world space destinations in the destination grid, or cannot reach that location from their initial positions and never travel there, or possibly because some combination of initial conditions have no valid observations of road users (invalid combination). Thus, in this example of 180 million cells, only 0.1 % of cells contain non-zero values. The sample size distributions for each initial condition modelled by the MPA for this dataset are shown in Figure 4.3 (one road user trajectory is one sample).

Different sparse arrays are tested to determine calculation performance. Early prototypes

³Sixty metres is the typical maximum size of the analysis zone used in this work and in similar studies.

of the discretized motion pattern motion prediction method made exclusive use of the Row-Based Linked List Sparse Matrix from the Python Numpy library, which supports flexible cell access (slicing) and efficient changes to sparsity (a trait that is useful for MPA training, as the size and complexity changes as more road user trajectories are added to the MPA). However, Row-based linked list sparse matrices are inefficient for arithmetic operation performance during motion prediction calculations, which are the most frequent operation. The Compressed Sparse Row Matrix, also from the Python Numpy library, does not have these same arithmetic operation limitations and is therefore used during motion prediction, instead of Row-Based Linked List Sparse Matrix (Jones et al., 2001–). Conversion of MPA from one sparsity structure to another is a trivial task.

Collision course calculation for discretized motion pattern motion prediction implements special rules given the discrete, probabilistic nature of the motion being modelled; discretized motion pattern motion prediction does not generate trajectories, instead, a probability map is generated. For a user pair at a single interaction instant at a time t_0 , probabilistic future positions for both road users A and B are predicted for instantaneous projected time $t_i > t_0$ according to MPA initial conditions. CPs are calculated by taking the joint probability of motion of both road users. These CPs are effectively MPA cells using the cell centroid as world space coordinates. Those cells where both A and B have a non-zero probability of arriving simultaneously generate discretized CPs:

$$P(AB_{x,y,t_0,t_i}) = P(A_{x,y,t_0,t_i} \cap B_{x,y,t_0,t_i}) = P(A_{x,y,t_0,t_i}) \times P(B_{x,y,t_0,t_i}) \quad (4.13)$$

where $P(A_{x,y,t_0,t_i})$ (resp. $P(B_{x,y,t_0,t_i})$) is the probability of the event of an arrival of A (resp. B) at position x, y at time t_i , given the initial conditions, and assuming events A_{x,y,t_0,t_i} and B_{x,y,t_0,t_i} are independent. This operation is illustrated in Figure 4.4.

Since this is a predicted collision, time $t_i - t_0$ serves as the TTC value for each CP at that time. However, multiple CPs are usually calculated and returned by the MPA for future times t_i . Each has a different probability. This is illustrated in Figure 4.5 showing a number of CPs all at the same coordinates x, y calculated for various t_i . Block intensity maps $P(CP)$, i.e. $P(AB_{x,y,t_0,t_i})$, which is usually very low. These values will be aggregated next, to make these values more accessible for further study of road design.

The aggregated probability of a predicted collision for road users A and B at interaction instant t_0 is the probability of the union of all predicted collisions over the time horizon, and at all significant points of collision:

$$P(AB_{t_0}) = P\left(\bigcup_{x=x_0}^{x_h} \bigcup_{y=y_0}^{y_h} \bigcup_{t=t_0}^{t_h} AB_{x,y,t}\right) \quad (4.14)$$

If collision events resulting from CPs are assumed to be mutually exclusive, then this calculation resolves into a simple sum of the probabilities of the individual CPs. Otherwise, calculation of the probability of the union of all collision events must follow the counting technique of the inclusion-exclusion principle (Mazur, 2009) for n events E_i , i.e.:

$$P\left(\bigcup_{i=1}^n E_i\right) = \sum_{k=1}^n (-1)^{k+1} \left(\sum_{1 \leq i_1 < \dots < i_k \leq n} P(E_{i_1} \cap \dots \cap E_{i_k}) \right) \quad (4.15)$$

However, with a time horizon of hundreds of interaction instants, this calculation becomes exponentially expensive and must be estimated instead. The union operation is estimated with an upper bound using Boole's inequality, which can handle many small probabilities well (e.g. $\eta < 0.05$)⁴, as their intersections are usually negligible in practice.

$$\begin{aligned} P\left(\bigcup_{x,y,t} AB_{x,y,t}\right) &= P\left(\bigcup_{(x,y,t) \text{ such that } P(AB_{x,y,t}) \leq \eta} AB_{x,y,t}\right) + P\left(\bigcup_{(x,y,t) \text{ such that } P(AB_{x,y,t}) > \eta} AB_{x,y,t}\right) \\ &\approx \sum_{x=x_0}^{x_h} \sum_{y=y_0}^{y_h} \sum_{t=t_0}^{t_h} P(AB_{x,y,t}) [P(AB_{x,y,t}) \leq \eta] + P\left(\bigcup_{(x,y,t) \text{ such that } P(AB_{x,y,t}) > \eta} AB_{x,y,t}\right) \end{aligned} \quad (4.16)$$

When probabilities are larger (larger than $\eta = 0.05$), Bonferroni's inequalities (Bonferroni, 1936) are used instead to approximate the $P\left(\bigcup_{(x,y,t) \text{ such that } P(AB_{x,y,t}) > \eta} AB_{x,y,t}\right)$ term, i.e.:

$$P\left(\bigcup_{(x,y,t) \text{ such that } P(AB_{x,y,t}) > \eta} AB_{x,y,t_i}\right) \leq \sum_j^{k_{odd}} (-1)^{j-1} S_j \quad (4.17)$$

$$P\left(\bigcup_{(x,y,t) \text{ such that } P(AB_{x,y,t}) > \eta} AB_{x,y,t_i}\right) \geq \sum_j^{k_{even}} (-1)^{j-1} S_j \quad (4.18)$$

for $k_{odd} = 2n + 1$, $k_{even} = 2n$, and $n \in \mathbb{N}$, with

$$S_k := \sum_{1 \leq x_1 < \dots < x_k \leq m} \sum_{1 \leq y_1 < \dots < y_k \leq n} \sum_{1 \leq t_1 < \dots < t_k \leq n} P(AB_{x_1,y_1,t_1} \cap \dots \cap AB_{x_k,y_k,t_k}) \quad (4.19)$$

⁴Note the use of Iverson bracket notation here to perform a conditional sum.

Once the two Bonferroni's bounds are calculated (to a depth k that can be supported by the processor, typically to about $k \leq 7$ iterations) the probability is approximated as the average of the two bounds. To simplify the calculation further, negligible $P(CP) < pThreshold = 0.00001$ is ignored entirely. Sensitivity to this parameter is evaluated in a later section. Note that $P(\bigcup_{x,y,t} AB_{x,y,t})$ is not a final probability of collision, as it assumes that road users do not react to the collision course. This component, $P(evasion)$, should be measured using the appropriate safety indicator, e.g. TTC, and any other relevant factors.

Calculation of the probability of the intersection of all collision events when using the counting technique of the inclusion-exclusion principle is important under the condition that a road user may be predicted to arrive at multiple CP simultaneously in the future. This is relevant for motion prediction methods that predict the partial positions of road users, for example: Point Set Prediction. Under these conditions, the probability of the intersection of collision events is non-zero only for those events that each model a separate portion of a road user. The mechanisms for implementing collision prediction of partial road user positions is left as future work.

4.3.3 Motion Prediction Methods Summary

Three methods for motion prediction, summarized in Table 4.1, are retained for further study as candidate models in the case analyses presented in Chapters 5 and 6: constant velocity, normal adaptation, and discretized motion pattern motion prediction. The former two form a baseline of methods in the literature while the latter is purpose-built to meet the requirements of modelling road user behaviour in non-linear movement tasks, e.g. in a roundabout.

Table 4.1 Summary of Prediction Methods Used

Description	Abbrev.	Model	Advantages	Disadvantages
Constant velocity	cv1	Context-free	Simple in concept and implementation	Cannot predict complex or non-linear motion
Normal adaptation	nad	Context-free	More forgiving of normal drift than cv1	Calibration of behavioural parameters
Discretized motion pattern	cmp	Empirical	Models naturalistic behaviour	Sampling issues and arbitrary discretization of space

Table 4.2 summarizes the different parameters (or initial conditions) used by the three motion

prediction methods. Generally speaking, any behavioural effects from parameters that are not taken into account by the motion prediction model have to be picked up empirically in an evasion component, $P(evasion)$, or even in surrogate safety method conversion factors.

Table 4.2 Summary of Collision Course and Evasion Parameters

Parameter	Parameter implemented for prediction model:			Future Implementation possible?		
	cvl	nad	cmp	cvl	nad	cmp
Initial Speed	Yes	Yes	Yes			
Context-adjusted speed	No	No	Yes	No	No	
Initial Heading	Yes	Yes	Yes			
Context-adjusted heading	No	No	Yes	No	No	
Random speed correction	No	Yes	Implicit	nad		
Random heading correction	No	Yes	Implicit	nad		
Non-linear motion	No	No	Yes	No	No	
Naturalistic behaviour	No	No	Yes	No	No	
Modelling desired destination	No	No	No	No	No	Yes
Visibility angle between users	No	No	No	No	Yes	Yes
Road user classification	No	No	No	No	Yes	Yes
Dynamic visual obstructions (large vehicles)	No	No	No	No	No	Feasible
Road user or vehicle state	No	No	No	No	No	No

The goal of surrogate safety methods is to perform analysis from externally observable information only. This means that data on individual characteristics of road users (e.g. driver impairment, state of distraction, vehicle mechanical condition) are not available and cannot be modelled, regardless of the model. This goal is justified on the basis that road design must be made for all road users indiscriminately, and thus must consider the needs of all road users in aggregate. Arguably, the population of road user can vary from one site to the next (e.g. in front of a school). Furthermore, collecting personal data of road users is ethically and logistically challenging. So, in very specific use cases, there may be a justification for introducing internal information to the model, but in general such information cannot or should not be included.

4.4 Classification of User Pairs

Classification of road user interactions pre-dates surrogate safety methods and originally appeared in the TCT. In the updated surrogate safety method, user pairs are categorized by relative position of road users, specifically angle of incidence Φ :

$$\Phi' = \cos^{-1} \left(\frac{(\vec{v}_1 \cdot \vec{v}_2)}{(\|\vec{v}_1\| \|\vec{v}_2\|)} \right) \quad (4.20)$$

$$\Phi = \begin{cases} 360^\circ - \Phi' & \text{if } \Phi' \geq 180^\circ \\ \Phi' & \text{otherwise} \end{cases} \quad (4.21)$$

between road users at any particular interaction instant using velocity vectors \vec{v}_1 and \vec{v}_2 of each road user. The orientation of road users and angle of incidence of collision between road users has implications regarding impact location on the colliding bodies and also road user reaction time and collision probability, given that visual perception of dangers is an operation that occurs within a limited field of view. Note that in practice the velocity vectors are not always perfectly aligned with the body of the road user. However, this effect is usually minimal, and can be further addressed operationally through good sources of data. Four main categories have been proposed throughout the literature (Elvik et al., 2009; Saunier et al., 2010; St-Aubin et al., 2013b), including :

- **Diverging interactions** are a type of situation between two road users whose predicted trajectories do not overlap. These are not collision courses at all, nor do they carry any significant real collision probability, as both road users move away from each other.
- **Head-on conflict** is a type of user pair categorization where both road users face each other (where the angle of incidence may be considered to be between 150° and 180°).⁵ This type of situation is common within opposing flows and left-turn conflicts at intersections, as well as on two-lane highways, but is otherwise rare.
- **Rear-end conflict** is a type of user pair categorization where both road users follow each other (where the angle of incidence may be considered to be less than 30°). This type of situation is common virtually everywhere, as road users are expected to travel in the same direction in all simple traffic corridors.

⁵This and the following angular definitions are arbitrary; there are no consistent conventions in the literature regarding conflict classification by angle.

- **Side-swipe conflict** is a type of user pair categorization where both road users travel next to each other without meeting the criteria above (where the angle of incidence may be considered to be between 30° and 150°). This type of situation is expected in multi-lane facilities of all types, especially at or near junctions. In fact this category may be subdivided into distinct groups based on lane-changing and merging actions.

This method of classification might be used in the future to calibrate the $P(\text{evasion})$ predictive power of specific SSM models or even motion prediction methods. For example, one would expect that road users facing each other in a head-on conflict are more likely to see each other and initiate evasive action than in a rear-end conflict situation, assuming identical TTC, and thus available time for perception, decision making, and evasive action, in either case. Alternative classifications might be made by manoeuvre type:

- **Car following** (Brackstone and McDonald, 1999) (and pedestrian following) is the most basic manoeuvre and it occurs virtually everywhere. This manoeuvre frequently leads to a collision course when the following road user accelerates to a faster speed than the leading road user. It is chiefly responsible for rear-end conflicts, which is why these conflicts are so predominant throughout all types of road geometry.
- A **converging manoeuvre** is a type of road user interaction where road users are forced to interact with each other due to the road configuration, often having reduced or no prior knowledge of each other's presence prior to trajectory negotiation. Thus, unlike in lane change manoeuvres or in following manoeuvres where collision courses may be a direct result of road users' intentions, in converging manoeuvres, collision courses are created by nature of the configuration of the roadway. In a converging manoeuvre, a gap acceptance procedure (Maze, 1981) is immediately triggered.

Rear-end conflicts and side-swipe conflicts are common in converging manoeuvres. Examples of converging manoeuvres include merging at highway on-ramps, and merging at any type of intersection with poor visibility, especially if the intersection is not signalized (or utilizes yield signs, as in the case of roundabouts).

- A **lane change manoeuvre** is a type of road user interaction where road users move into a different lane of a multi-lane facility. As with converging manoeuvres, if road users are already present in the destinating lane, then a special gap acceptance procedure (Maze, 1981) occurs to allow the road user to enter the lane without colliding with

any road users already present. Rear-end conflicts and side-swipe conflicts are common in lane change manoeuvres.

- Much like with lane change manoeuvres, a **crossing manoeuvre** is a type of road user interaction where one road user attempts to cross a conflicting stream of traffic. This process is distinct from converging manoeuvres in that it typically involves orthogonal approaches as opposed to tangential approaches. This usually carries greater risks, as awareness or visibility may be further reduced and the differential velocity between road users is greater. Examples include pedestrian-car interactions or road users crossing dense traffic streams. This risk is typically compensated for by giving road users a chance, and often explicit instructs, to come to a complete stop before proceeding through a conflicting stream of traffic. Nevertheless, these manoeuvres may involve head-on conflicts and side-swipe conflicts with road users in the crossing stream, and can also result in rear-end conflicts with road users in the originating stream. Examples of crossing manoeuvres include any manoeuvres through traditional signalized and sometimes unsignalized intersections, especially involving right and left turns.

These interaction classifications convey particular information about collision course timing: in a converging manoeuvre, less time is available to, or used by, the road user to successfully execute a gap acceptance, whereas in a crossing manoeuvre, road users may take more time to execute a gap acceptance manoeuvre as they will likely be at a standstill. When performing a gap acceptance manoeuvre is time-sensitive, safe following distances and safety margins may have to be compromised, leading to more vehicular proximity overall, and more collision courses, especially if operating speeds differ.

4.5 Surrogate Safety Measure Aggregation Methods

Some SSMs are generated at every interaction instant. This is the case for most collision-course-based SSMs, especially safety indicators such as TTC, and speed as a SSM, but not for others, such as PETs.

The series of safety indicators generated in this way produces a time series of safety indicators, as illustrated for all three motion prediction methods for the same user pair in Figure 4.6. As with Figure 4.2, the expected evolution of TTC is plotted. Evasive action, intentional or not, is recognizable at around $t_0 = 1.5$. For those motion prediction methods which predict multiple CPs at each interaction instant, TTC measures for each CP are plotted individually, with plot intensity corresponding to normalised probability of each CP. The solid line is a weighted average of these probabilities. It should be noted that some motion

prediction methods produce many more safety indicators than others. Also, safety indicators are somewhat noisy between interaction instants; this is partially due to tracking errors and inaccuracy, and also due to aliasing between trajectories the MPA.

Analysing these measures in a disaggregated manner (sampling all interaction instants) can cause interactions with more interaction instants and bias the results of some sites more than others for a number of reasons:

- Sites with trajectory data from cameras with higher frame rates oversample individual interaction instant observations compared with sites with trajectory data from cameras with lower frame rates. This is relatively easy to correct.
- Some bias can be introduced through low tracking quality due to unknown or difficult-to-identify factors, such as video quality (different cameras, unclean optics, poor lighting conditions) resulting in shorter or lower quality trajectories, and possibly impacting the number of observations.
- Larger analysis zones produce more observations at a single site, since road users dwell within a scene for a longer period of time given the same speed.
- Road users travelling more slowly also cause an oversampling of interaction instant observations.

Arguably, it is desirable to keep the last factor in a road safety model: theoretically, more exposure *should* result in increased probabilities of collision (at a site, i.e. the expected number of collisions) in the same way that driving further distances increases lifetime collision probability and in the same way that throwing more successive dice increases the probability of obtaining at least one particular roll. However, no historical accident data exist yet to construct such a quantitative collision probability calibration. In any case, in order to solve these aforementioned issues, interaction instants are aggregated to produce aggregated SSMS per user pair, often representing the a “worst-case” or “critical” interaction instant.

Several aggregation methods have been proposed and presented (in St-Aubin et al., 2015b). These select a single value to represent the user pair interaction, and include:

- **Disaggregated:** all safety indicators are kept. This method is not recommended for reasons outlined earlier, though it may be revisited in future with a more sophisticated adaptation of the surrogate safety method.

- **Minimum value:** for all safety indicators of a user pair, the lowest observed value is kept, if the relationship between the safety indicators and some aspect of road safety is positive and monotonic; otherwise, the highest is used. This concept is based on common usage, throughout the literature, of the minimum TTC, TTC_{min} (e.g. Laureshyn, 2010), and is related to the safety indicator time series presented in Figure 4.2 (see also Hayward, 1971). One notable problem with this approach is that observations tend to be noisy at the interaction instant level, causing outlier bias when selecting extreme values.
- **15th centile:** this is identical to minimum value aggregation, however the 15th centile is used instead of the 0th centile, to be more statistically robust to noisy data, if the relationship between the safety indicators and some aspect of road safety is positive, otherwise, the 85th centile is used.
- **Maximum probability:** for safety indicators with non-deterministic collision course probabilities, such as CPs predicted from discretized motion pattern motion prediction, the safety indicator associated with the CP with greatest probability is used to represent the user pair.

These methods and their technical definitions applied to TTC safety indicators are summarized in Table 4.3.

Table 4.3 Summary of Aggregation Methods (TTC)

Description	Abbrev.	Values per user pair	Measure Criterion	Measure Centile
Disaggregated	a11	all values	None	N/A
Unique minimum pair	min	unique value	Safety indicator	0.0
Unique 15 th centile pair	15t	unique value	Safety indicator	15.0
Unique maximum probability	maxProb	unique value	Predicted collision probability	100.0

4.6 Evaluating Road Safety Using Surrogate Safety Measures

This section describes analysis strategies available for SSMs and presents an experimental analysis of the different TTC motion prediction methods and aggregation methods. The experimentation is performed on sample roundabout merging zones in the province of Québec.

This data is discussed in much more detail in the next chapter as part of a full road safety analysis case study, specifically section 5.2. In summary, this data consists of 196,808 road users at 35 merging zones across a wide variety of land uses, lane configurations, speed limits, and various geometric factors.

4.6.1 Choice of Safety Evaluation

Two general approaches in the literature for evaluating road safety based on SSMs are followed: safety continuum comparison (SCC) and serious event comparison (SEC). Each approach has several variants.

Serious Event Comparison

In a **serious event comparison (SEC)**, the significance of SSM on road safety is evaluated from the start, treating all and only events (e.g. user pairs) with values meeting a threshold value ζ as contributing to some aspect of road safety. These are termed “serious events.” Correspondingly, this approach tends to treat events as binary units of collision exposure: the event is either classified as dangerous or it is not.

The approach has been traditionally used in the TCT where road user interactions are classified in a similar binary manner as being either (serious) traffic conflicts or not. This approach has been repurposed in modern surrogate safety method by applying threshold values to SSMs. TTC, for example, is typically classified as a dangerous event if it has a value below $\zeta = 1.5seconds$, corresponding to commonly-cited average human reaction times used in literature (Hydén, 1987; Green, 2000). New and old surrogate safety methods then attempt to predict expected numbers of collisions using traffic conflict conversion factors (e.g. Hydén, 1987).

Normally, the definition of a traffic *event* is conducted in accordance with the desired aggregation method (see Table 4.3). Then, the number of serious events SE is counted among those SSMs which meet the threshold ζ criterion.

Given the collision course uncertainty modelled in the new surrogate safety method, a modified, modernized SEC methodology is warranted on the basis that not all collision courses will affect road safety (e.g. collision probability) equally. Therefore, a new methodology is proposed: the weighted SEC. The expected number of serious events wSE is calculated by summing all events which meet the desired safety indicator *ind* threshold and weighing each by its associated $P(CP)$ (depending on aggregation method), such that

$$wSE = \sum ind \times P(CP)[ind < \zeta] \quad (4.22)$$

or

$$wSE = \sum ind \times P(CP)[ind > \zeta] \quad (4.23)$$

Alternatively, $P(CP)$ may be substituted with $P(\text{collision})$ from equation 4.8 to account for any $P(\text{evasion})$. Note that weighted SEC only differs from classic SEC if predicted collision courses make use of a calculated $P(CP_e)$. Context-free motion prediction methods which do not have collision courses probabilities—thus where a value of $P(CP_e) = 1$ is assumed—are not affected by the weighing process and yield results identical to classic SEC.

Safety Continuum Comparison

In a safety continuum comparison (SCC), *all* SSMs, regardless of their value, are used for evaluation. In this approach, it is supposed that each SSM contributes to collision probability (e.g. how it affects $P(\text{evasion})$) and collision severity, but only in proportion to its value, and possibly in a non-linear relationship. For example, a collision course with a TTC value predicted 5 s into the future is not ignored outright as a safety concern, however its contribution is diminished because of the low reliability of such a prediction, i.e. the combined effect of $P(CP)$ and $P(\text{evasion})$. Similarly, speeds in excess of 5 km/h over the design speed are of less concern than more egregious instances of speeding, but are still not necessarily entirely negligible factors of road safety.

SCC makes fewer assumptions than SEC, but the model it relies on is partially incomplete. Predicting expected numbers of collisions from SCC requires a calibrated $P(\text{evasion})$. On the other hand, SEC does not attempt to calibrate $P(\text{evasion})$ at all. In any case, in the meantime, a complementary qualitative analysis using mean SSMs or SSM distribution is possible. This comparison is possible on the basis that $P(\text{evasion})$ for the same SSM is constant between populations of road users and that exposure is controlled for. One such comparison for TTC SSMs is illustrated in Figure 4.7.

A right-shift in a TTC cumulative probability distribution suggests improved road safety, as a higher proportion of events has a higher TTC, providing more time for a larger share of road users to evade collisions. Thus, a right-shift of “low-value” “high-risk” TTC cumulative probability and a left-shift of “high-value” “low-risk” TTC cumulative probability could be

suggestive of improved safety (Figures 4.7a and 4.7b). However, such a comparison may not be conclusive if a smaller proportion of “low-value” “high-risk” and “high-value” “low-risk” TTC events is obtained in exchange for an increase in “medium-risk” TTC events, regardless of magnitude (Figures 4.7c and 4.7d). In the same way, comparison of a mean TTC can, at best, be suggestive of improved safety, but cannot be overall conclusive.

Meanwhile, a number of researchers have argued that an increase in “medium-risk” TTC events may be beneficial to road safety, as such events are serious enough to cause discomfort and solicit attention from road users, but not serious enough to cause an appreciable increase in collisions beyond the benefits of increased road user awareness (Svensson, 1998; Svensson and Hydén, 2006; Saunier et al., 2011b; Saunier and Mohamed, 2014). What constitutes “medium-risk” TTC is still up for debate.

Ultimately, both models make similar assertions about how individual SSMs impact collision probability or collision severity. The SCC approach is more complex, but also more generalizable and potentially much more nuanced and powerful, given a fully calibrated model. SEC is primarily limited by its blunt, binary nature, however it is simple to define and has had some validation performed using historical accident data in a single study (Hydén, 1987).

4.6.2 Comparison and Choice of TTC Measures

TTC-based measures generated from different prediction methods and aggregation methods have been briefly evaluated. Generally speaking, discretized motion pattern motion prediction and unique 15th centile aggregation is proposed as a superior SSM model to other methods on mostly intrinsic grounds: despite correctable issues with sampling and restricted time horizons, this method offers intrinsic benefits with more naturalistic motion prediction, explicit collision probability calculation, and robustness against noisy data (St-Aubin et al., 2015b). In this section, additional tests are performed on samples of test data to better quantify the use of the varying motion prediction and aggregation methodology for TTC measures.

Collision-Course Probability Threshold

Experimentally, it is shown that collision course probability thresholds for ignoring negligible values and simplifying calculations have a negligible effect on site ranking. Figure 4.8, which compares collision course probability thresholds 0.01 and 0.001, illustrates that site ranking is negligibly affected by choice of probability thresholds but that SEC is potentially sensitive

to choice of collision course probability threshold (St-Aubin et al., 2015b).

Time Series Evaluation

It should be remembered from Figure 4.2, and by definition of TTC, that TTC values are expected to decrease linearly at a rate of one second per second over the time series of each individual user pair, if the road users continue on their predicted trajectory. These time series are examined on an individual basis. Figure 4.9 plots the entirety of all TTC interaction instants for each user pair throughout the entirety of the test data for each motion prediction method, as well as the average evolution.

One problem with these figures becomes apparent: time series lengths differ from one user pair to the next, and, as was illustrated in Figure 4.6, the context-free motion prediction methods tend to generate significantly shorter time series. Time series lengths cannot be normalized without warping time series shape. Thus, a better option is to present this data by offsetting the entire time series so that the critical minimum TTC value, TTC_{min} , is always located at the same time (e.g. $t = 8s$). A more involved approach used elsewhere in the literature involves maximizing time series similarity via LCSS (Saunier and Mohamed, 2014).

The results demonstrate that, overall, the expected evolution of TTC linearly decreases at a rate of 1 s of TTC per 1 s of time. Deviation from this expected evolution is interpreted as control input, i.e. change introduced by road users likely reacting to the collision course itself. It is interesting to note that the average TTC_{min} measured by normal adaptation and discretized motion pattern motion prediction is just slightly higher than 2 s, while constant velocity motion prediction TTC_{min} is measured closer to the literature-standard average reaction time of 1.5 s. This might suggest that road users are actively anticipating collision courses, well before they become serious.

Given the curved trajectories, constant velocity may be suffering sampling issues. Constant velocity may also be overestimating serious collision courses at stop or yield which are present at all sites in the study, by not accounting for changing conditions of road user desired speed. It is difficult to draw conclusions from constant velocity, other than to say that it is problematic for this reason. Note that some stripes can be seen in some discretized motion pattern motion prediction time series. This is probably due to the discrete nature of discretized motion pattern.

Serious Event Comparison Evaluation

Using the SEC approach, using $\zeta = 1.5$ s as a threshold value for TTC (as prescribed in the literature (Green, 2000)), each of the 35 analysis zones are ranked in order of number of traffic events per hour below 1.5 s. Spearman's ranking correlation is performed across each method. Note that SEC automatically uses a time series value aggregation, so the method of using all disaggregated values is not included in this analysis.

Table 4.4 Spearman's Ranking Correlation for Serious Event Comparison

		cvl			nad			cmp	
		min	15t	maxProb	min	15t	maxProb	min	15t
cvl	15t	0.9926							
	maxProb	0.9867	0.9951						
nad	min	0.9955	0.9919	0.9868					
	15t	0.9811	0.9935	0.9916	0.9872				
	maxProb	0.9709	0.9851	0.9935	0.9729	0.9882			
cmp	min	0.7093	0.7091	0.7176	0.6995	0.6936	0.7152		
	15t	0.7090	0.7065	0.7153	0.6926	0.6881	0.7107	0.9894	
	maxProb	0.6688	0.6615	0.6672	0.6457	0.6343	0.6562	0.9348	0.9553

The results of this ranking suggest that discretized motion pattern motion prediction presented in Figure 4.4 yields results different from both constant velocity and normal adaptation motion prediction and that these latter two are essentially interchangeable. The results also suggest that the choice method of single value aggregation has little impact on site ranking, especially for constant velocity and normal adaptation motion prediction, but still has a slight impact on discretized motion pattern motion prediction. The results are summarized in Table 4.4.

Figure 4.10 plots an ordered ranking of sites based on hourly number of user pairs (events $i = 1..n$) with at least one TTC indicator below threshold value $SEC_{TTC < \zeta}$ (left vertical axis), e.g. following

$$\dot{SEC}_{TTC < \zeta} = \frac{\sum_i [TTC < \zeta]}{t} \quad (4.24)$$

versus the fraction of serious events ϕ (right vertical axis)

$$\phi = \frac{t \dot{SEC}_{TTC < \zeta}}{n} \quad (4.25)$$

The graph also includes a plot of the number of hourly user pairs with at least one instantaneous TTC measure below threshold value versus hourly number of user pairs, and demonstrates that the two are moderately correlated ($R^2 = 0.49$). This is expected by conventional knowledge in road safety theory: collision rates, or their surrogates, i.e. events with low TTC values, increase non-linearly with exposure (i.e. user pairs).

The distribution of serious events across sites in Figure 4.10 follows a classic Pareto distribution (20 % of sites account for 80 % of serious events). A small number of sites contain disproportionately large hourly proportions of serious events, however, given that the exposure (hundreds per hour) and number (single observations per hour) of serious events is so small, these are unreliable or outlying observations. The remaining sites with significant numbers of user pairs demonstrate a more stable relationship. As the number of user pairs increases, so does the total number of serious events, but the rate of serious events per user pair decreases (or, at the very least, it does not increase proportionately for the busiest of sites), likely explained by the “safety in numbers” effect.




Figure 4.3 Sample size statistics for initial conditions of an example motion pattern array. The top figures plot the number of initial condition combinations with corresponding sample sizes (i.e. to study the distribution of initial condition combinations with very few trajectory samples). The left figure is the same as the right with simply a shortened x-axis. The bottom figure plots the sample size for each combination of initial conditions ordered by size.




Figure 4.4 Probability map based on the joint arrival of road users for discretized-motion-pattern-based collision course prediction.

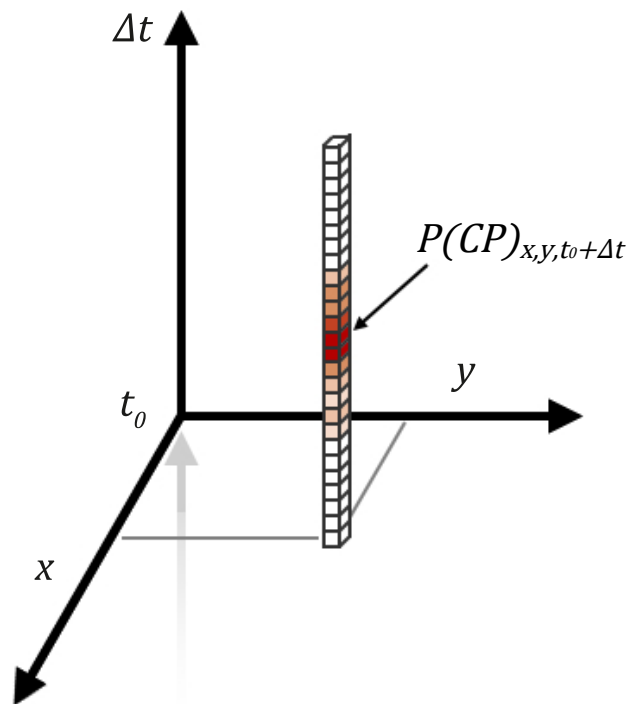


Figure 4.5 Collision points at a single position for a single collision course. They are predicted discretely at each cell in world-space of the MPA grid, at every time step over a given time horizon into the future. Probability of a collision point is highlighted (darker indicates a greater probability).




Figure 4.6 Time series of TTC indicators for a single user pair calculated using three different motion prediction methods. An interaction instant may have multiple TTC indicators (if predicted probabilistically) and these are shaded according to their probability of collision. Each curve represents the weighted average TTC observation at that interaction instant.

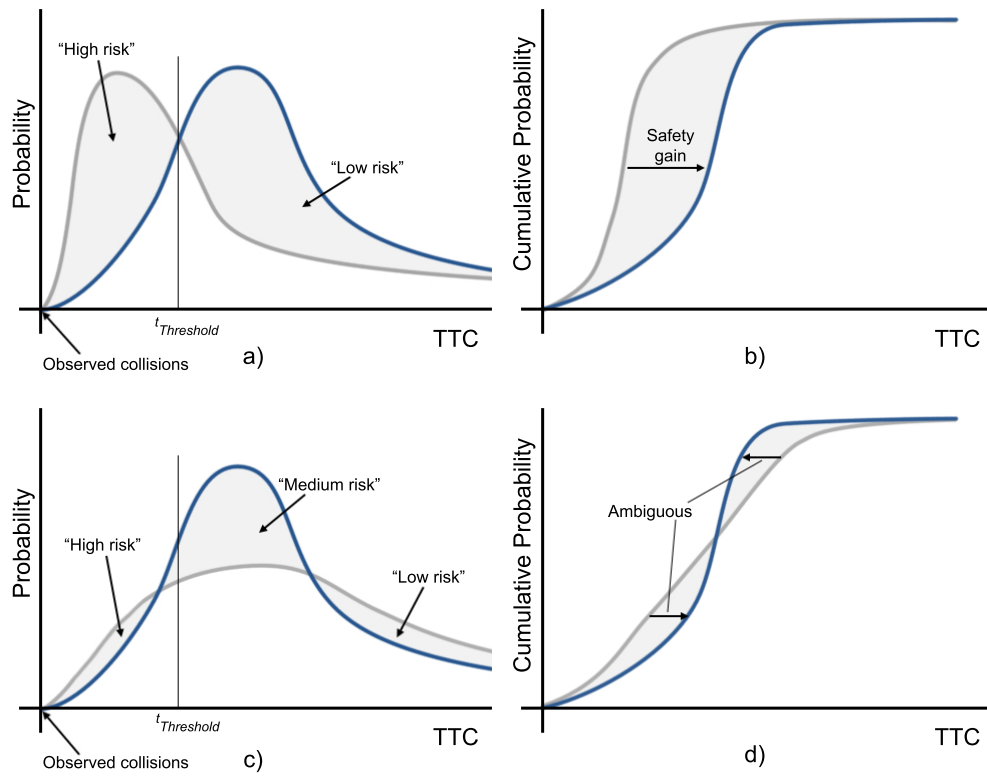


Figure 4.7 TTC distribution comparison for a before-after or cross-sectional safety continuum comparison.






Figure 4.8 Serious event comparison of hourly events for collision-course probability thresholds 0.01 and 0.001 across sites.



(a) Constant velocity

(b) Normal adaptation



(c) Discretized motion pattern

Figure 4.9 Disaggregated and aggregated time series of all user pairs in test data.




Figure 4.10 Comparison of hourly serious event counts (blue, ranking criteria) versus percentage of events rated as serious (orange). The inset graph shows a moderate correlation (0.4866) between hourly serious event counts and number of road users.

Weighted Serious Event Comparison Evaluation

For the same prediction method (in this case, only discretized motion pattern motion prediction, as it alone, of the test methods, makes use of dissimilar calculated probabilities of predicted collision) and aggregation method, the Pearson linear correlation between SEC site measures and weighted SEC site measures is found to be 0.90, while the Spearman rank correlation is 0.97. Thus the two methods of comparison produce substantially similar, but not identical results in a comparative manner. However, the total number of expected events is reduced in order to better reflect calculated probabilities of collision. This helps to counterbalance the tendency of discretized motion pattern motion prediction to predict many more low-probability CPs than constant velocity and normal adaptation. The safety indicator threshold used is identical to that of ordinary SEC.

Safety Continuum Evaluation

Using the SCC approach, distributions for different motion prediction and aggregation methods are compared in Figure 4.11. As noted in SEC, constant velocity and normal adaptation motion prediction yield visually similar results. These motion prediction methods also yield a large proportion of values suspiciously close to 0 s, which can be considered serious. With these TTC values and this many user pairs (176,749), many more collisions would be expected as part of the data (in reality, only one real collision was observed throughout the entire data set).

Also, similar conclusions can be drawn for aggregation methods: aggregation method has the greatest impact on conclusions for discretized motion pattern motion prediction. Naturally, using extreme values biases the results, though the difference between unique minimum pair and unique 15th centile pair is less pronounced than expected, after aggregation to the site-level. Overall, regardless of motion prediction, all distributions appear to be roughly gamma-like when aggregated to the site level. The shape of these distributions will be investigated further in the next chapter. Some distribution clusters exist, but for the most part, some sites stand out as being conclusively less safe than others via SCC after examination of shifts in distribution.




Figure 4.11 Cumulative distributions of TTC indicators for each site by aggregation method and prediction method. In red is the cumulative distribution of TTC indicators with sampling size labelled and opacity set to a percentage of the largest contributing sample. In black is the total distribution of all values across all sites.

Table 4.5 presents the Spearman rank correlation for mean TTC value observed at each site, across motion prediction and aggregation methods, much like in Table 4.4. This time, median TTC is used as a stand-in for the disaggregated method. Since probabilities for constant velocity and normal adaptation are constant, by design of these algorithms, maximum probability selects values in an unpredictable manner. Results for these two cases should be ignored. Conclusions are similar to those of SEC, though maximum probability discretized motion pattern motion prediction is found to be strikingly different from minimum pair discretized motion pattern motion prediction.

Table 4.5 Spearman's Ranking Correlation for Safety Continuum Comparison

Aggregation Method for Constant velocity			
	Median	Minimum pair	15 th centile pair
Median	1.0000		
Minimum pair	0.9271	1.0000	
15 th centile pair	0.9386	0.9933	1.0000
Aggregation Method for Normal adaptation			
	Median	Minimum pair	15 th centile pair
Median	1.0000		
Minimum pair	0.8758	1.0000	
15 th centile pair	0.8932	0.9820	1.0000
Aggregation Method for Discretized motion pattern			
	Median	Minimum pair	15 th centile pair
Median	1.0000		
Minimum pair	0.5911	1.0000	
15 th centile pair	0.7843	0.8489	1.0000
Maximum probability	0.8495	0.4713	0.6518

4.6.3 Summary of Surrogate Safety Measures

Table 4.6 summarizes the SSMS and associated parameters available for analysis. For the most part, these should serve as the dependent variables for use in a road safety analysis.

In addition to user pairs, more classic measures of exposure are generated and considered for study. These are listed in Table 4.7 and include exposure attached to individual traffic events, i.e. user pairs with at least one safety indicator $P(CP) > 0$, and exposure aggregated to the unit of analysis (i.e. at each analysis zone).

Table 4.6 Safety Indicators

Variable (user pairs)	Factor description	Type (units)
<code>ttc_val</code>	TTC value for corresponding motion prediction and aggregation method (e.g. <code>15t</code> , <code>cmp</code>)	Numerical (seconds)
<code>ttc_prob</code>	Collision course probability	Numerical (0-1)
<code>lower_inst_speed</code>	Smaller of two instantaneous speeds of a user pair at the instant of prediction of a collision	Numerical (km/h)
<code>higher_inst_speed</code>	Larger of two instantaneous speeds of a user pair at the instant of prediction of a collision	Numerical (km/h)
<code>interinst_angle</code>	Angle of incidence Φ between both road users of a user pair (see equation 4.20)	Numerical ($^{\circ}$)
Variable (aggregated to the unit of analysis)	Factor description	Type (units)
<code>mean_speed_ph</code>	Mean speed of all road users within the unit of analysis	Numerical (km/h)
<code>mean_ttc_X</code>	Mean TTC of all user pairs for corresponding motion prediction and aggregation method (e.g. <code>15t</code> , <code>cmp</code>)	Numerical (seconds)
<code>ttc_count_15_weighted</code>	Weighted hourly expected number of serious events wSE for corresponding motion prediction and aggregation method (e.g. <code>15t</code> , <code>cmp</code>)	Numerical (expected events/h)

4.7 Summary

In this chapter, the basic road safety analysis framework is extended to predict *expected collision p.a.* from modernized and systematically measured and defined SSMs following the proposed surrogate safety methodology. A model of probability of evasion is also introduced (though not evaluated—this will have to be formally tested as part of future work with access to better data) following a review of theoretical concepts of typical collision mechanics, such as reaction time, distractions, visual obstructions, etc. The core contributions are the formal definition of user pairs within the analysis zone, the generalized definition of the collision course (re-purposing the crossing zones and CPs concepts), along with an associated collision probability space and aggregation methods. A specialized motion prediction method, i.e. a discretized motion pattern which is capable of handling non-linear motion environments and

Table 4.7 Measures of Exposure or Related to Exposure

Variable (user pairs)	Factor description	Type (Units)
hour	Hour of the day	Numerical (0-23)
five_minute_exposure	Number of road users present in the analysis zone 2.5 minutes before and after any instant, e.g. a collision course	Numerical (veh/5 min)
two_minute_exposure	Number of road users present in the analysis zone one minute before and after an event	Numerical (veh/2 min)
one_minute_exposure	Number of road users present in the analysis zone 30 s before and after an event	Numerical (veh/min)
fifteen_second_exposure	Number of road users present in the analysis zone 7.5 s before and after an event	Numerical (veh/15 min)
Variable (aggregated to the unit of analysis)	Factor description	Type (Units)
user_pairs_ph	Hourly number of road user interactions (non-instantaneous) within the analysis zone	Numerical (int/h)
inflow_ph	Hourly flow rate within the unit of analysis	Numerical (veh/h)
inflow_phpl	Hourly flow rate within the unit of analysis divided by the number of inflow lanes	Numerical (veh/h)

naturalistic behaviour efficiently, is also innovated.

Various SSMs are presented, and three general-purpose SSMs are selected for application in the case study of roundabouts to follow. These are instantaneous measures of speed, yPET, and TTC, from discretized motion pattern motion prediction aggregated to the 15th centile per user pair.

Overall, better historical accident data will be needed as part of a future project to formally test collision probability from SSMs, such as TTC, and to calibrate a proposed $P(\text{evasion})$. In the meantime, major issues with subjectivity and reproducibility (Williams, 1981; Krusysse, 1991; Chin and Quek, 1997) are addressed thoroughly in this chapter, especially with the introduction of more sophisticated and context-aware motion prediction methods (e.g. discretized motion patterns) which could potentially improve TTC predictive power as an SSM, or even remove dependency on intermediary SSMs entirely by modelling collision probabil-

ity directly. Furthermore, a number of user pair classifications are proposed which may be hypothesized to affect $P(evasion)$.

The road safety analysis strategy SCC was introduced and tested alongside SEC. Meanwhile, to make SEC methodology more viable, it has been argued that more accurate evasive action time modelling will be necessary.

CHAPTER 5 CASE STUDY OF QUÉBEC ROUNDABOUTS

This chapter guides the reader through a full application of the video-based traffic-data collection system presented in Chapter 3 in a typical traffic engineering and road safety study. For the purposes of road safety analysis, the surrogate safety methodology presented in Chapter 4 is applied to the trajectory data obtained from the video-based traffic-data collection system. Specific SSMs used for the road safety analysis include speed, yielding post-encroachment time (yPET), and time-to-collision (TTC) using discretized motion patterns and 15th centile aggregation.

In this chapter, a road safety analysis is performed in the form of a cross-sectional study of road user behaviour on a sample of roundabouts in the province of Québec. The objective is to explain aspects of road safety via surrogate safety methods from specific elements of geometry and land use parameters. As stated earlier, roundabouts are ideally purposed for an application of surrogate safety methods given that

- roundabouts feature merging zones where traffic interactions are governed by road user behaviour exclusively;
- these interactions are complex in nature, requiring a deeper understanding of the underlying mechanisms driving them; and
- roundabouts have relatively low traffic-volumes (and collisions), resulting in sparse historical accident data.

Meanwhile, a surrogate-safety-method-based study of roundabouts is justified on the grounds that these studies are relatively uncommon in the literature (especially for North American roundabouts) and that this research was conducted as part of a larger research making a road safety analysis of roundabouts in Québec, providing an opportunity and incentive to undertake data collection for this project.

5.1 Roundabout Design Elements

Roundabouts have been promoted as a safer alternative to traditional intersections, promising reductions in traffic conflicts types, (and more specifically, “fewer conflict points” Rodegerdts et al., 2010). Studies have historically demonstrated reductions in serious vehicular collisions (Hydén and Várhelyi, 2000; Persaud et al., 2001; Gross et al., 2013; Jensen, 2013), and

reductions in speed (Hydén and Várhelyi, 2000; St-Aubin et al., 2013a). However, issues with vulnerable users, such as pedestrians and cyclists, have been identified in the literature (Hydén and Várhelyi, 2000; Daniels et al., 2010; Cumming, 2012), particularly with respect to disabilities such as visual impairment. Meanwhile, few, if any, studies have attempted to examine the underlying collision mechanics at a microscopic level, except for possibly H. Sadeq (Sadeq, 2013), who examined a single roundabout extensively using automated trajectory data collection and a TCT-inspired, i.e. leaning more towards use of SEC, analysis. In the following sections, the specific design of the roundabout is dissected to lay the groundwork for a microscopic analysis of road user behaviour and aspects of road safety.

5.1.1 Merging Zone as a Unit of Analysis

At the macroscopic level, roundabouts operate similarly to ordinary traffic intersections serving two or more traffic corridors. However, roundabouts are distinguished from traditional four-way intersections at the microscopic level in that, instead of mixing all traffic movements from all approaches within the same space by alternating the right of way of conflicting movements, roundabouts isolate conflicting movements into separate merging zones, effectively removing left-turn manoeuvres (road users proceed around the ring if they wish to exit by the left-most exit) and head-on conflicts. This isolation effect is not only physical (Rodegerdts et al., 2010); it has also been highlighted in the literature as being psychological, with observations that the height of the central island contributes positively to a reduction in the number of observed collisions (collision probability) by removing visual distractions and by forcing the road users to pay attention when approaching the intersection (Jensen, 2014).

In this way, movements of each roundabout approach are isolated via dedicated merging zones. If one were to cut the ring of a roundabout on one of its sides and unroll it, one would find a shape that greatly resembles a one-way stretch of limited-access highway: a section of closed track with a series of approaches and exits. A typical roundabout merging zone will have one approach, followed by one exit located next to the counter-clockwise adjacent approach (in right-driving regions). This approach/exit combination contains a merging zone not unlike that of a merging zone found at limited-access highway junctions and ramps where routing decisions, manoeuvres, and merges—along with side-swipe conflicts and, to a lesser degree, rear-end conflicts—are forced to occur. Thus a merging zone exists for each approach around the ring. In an ideal, symmetrical, four-legged roundabout (Figure 5.1), each of these zones corresponds to one quarter of the roundabout ring, thus the merging zone can effectively also be thought of as a *quadrant*. However, in practice, variability in the approach

angle, directions of travel, number of lanes, and number of branches often yields more or less than four merging zones of variable size at a single roundabout.




Figure 5.1 Typical four-way roundabout sectioned into component merging zones.

Given the purpose of this study, that is, to explain aspects of road safety from elements of geometric design and land use at roundabouts, an inventory of roundabout design elements constituting potential contributing factors to road safety must be produced. As such, merging zones are used as individual units of analysis, instead of roundabouts as a whole. This choice was made for a number of reasons including:

- the scope of analysis of individual merging zones is more relevant to investigating individual collision mechanisms than merging zones in the aggregate, i.e. the roundabout as a whole;
- roundabouts are mostly symmetrical, incorporating a sequence of merging zones in a ring, with one merging zone for every roundabout approach;
- despite apparent symmetry, a great deal of variability in potential contributing factors can nevertheless be found experimentally from one merging zone to the next, and these differences are easily parametrized;

- road user interactions are generally isolated between the two ends of the merging zone; and
- at the start of the project, obtaining and installing 360° field of view cameras was deemed too logistically challenging.

This is in contrast to historical accident analysis methods which evaluate road safety at intersections—including roundabouts—as a whole, as the historical accident data collected seldom includes location data precise enough to geolocate collisions at a smaller unit of analysis.

The variability of geometric design factors between individual merging zones can be parametrized to build a database of merging zones using more than 60 variables corresponding to road geometry, signage, markings, and environment, although as will be seen later, many of these factors are collinear, or of lesser importance.

Figure 5.2 illustrates the location and types of typical vehicle-vehicle traffic conflicts in a typical roundabout merging zone. At all times, in any type of road infrastructure, for that matter, vehicles may be subject to rear-end conflicts. Meanwhile, merging zones produce converging manoeuvres which are characterized by forced road user trajectory negotiations with new road users over a short distance (i.e. the merging zone), and which may therefore generate additional side-swipe conflicts. Finally, in multi-lane facilities, interaction complexity increases further as the number of lanes and any required lane changes increases, not only generating additional side-swipe conflicts from ordinary lane change manoeuvres, but also further increasing the complexity of converging manoeuvres. This is why multi-lane roundabouts are so problematic (Rodegerdts et al., 2010): side-swipe conflicts are compounded by two unique and separate traffic processes. The distinction between these two processes is that, in converging manoeuvres, road users are forced against each other at specific locations of the road due to mandatory lane changes, while in lane change manoeuvres, road users make discretionary lane changes. Thus, it is the merging zone that is of greatest interest for investigation of collision mechanisms, as it theoretically generates the most traffic conflicts and traffic flow friction.

One drawback with the use of merging zones as a unit of analysis is that it somewhat limits comparative analysis of roundabouts having more traditional intersections, since road user interactions at traditional intersections are not isolated in the same manner as at roundabouts. This issue is of little consequence for this chapter, as no traditional intersections are under investigation as part of the study, but future comparison between results for these roundabouts and for traditional intersections could be challenging. A comparison would re-




Figure 5.2 Vehicle-vehicle conflict zones for a typical roundabout merging zone.

quire the complete examination of all sections of both the roundabouts and the intersections being compared. Another limitation of this approach is that road user interactions *between* merging zones are not studied. In practice, this is not exactly true, since analysis zones must be constructed to cover the sections of road leading into the merging zone and effectively also leading out of the merging zone. Furthermore, the only user pairs which cannot be generated in this way are those which originate in significantly different merging zones and which never enter any merging zones, or approaches and exits adjacent thereto, simultaneously. Arguably, these road users are not in conflict with each other, as supported by evidence (in Jensen, 2014).

5.1.2 High-Level Interpretation Module: Roundabouts

The automated video-based traffic data collection and analysis system, and surrogate safety method framework presented in Chapters 3 and 4 respectively, are general frameworks for traffic analysis at any type of road installation and for any type of road user. To perform

roundabout-specific measurements, i.e. measurements taken with respect to the geometry or for custom analysis, HLI modularity is employed and illustrated in this subsection.¹ More specifically, this module stores roundabout metadata, then uses this metadata to calculate measures specific to the geometry of the roundabout: yPET, movement profiles, and flow ratios.

Metadata

Using HLI modularity, the roundabout, approach and exit lane numbers, as well as the respective location of their intersection with the merging zone, are tagged. Road user are deemed to have entered or exited the merging zone when their curvilinear position—recall that curvilinear positions are bound to lane alignments—crosses one of these boundary locations, i.e. crosses the roundabout conflicting lane, where the conflicting flow circulates, or the approach edge, and hence the merging zone. This is illustrated in Figure 5.3. With this metadata, the HLI-specific calculation can be made and compiled for geometric-specific measures, notably yPET.

Yielding Post-Encroachment Time

In the previous chapter, TTC was introduced as a general-purpose SSM using motion prediction. While collision courses are the basis of continued analysis of roundabout merging zones, PET offers additional insight into road user behaviour free from the necessity of modelling collision courses via motion prediction. It does so by measuring the proximity of observed trajectories that overlap but do not collide—they *encroach* (Allen et al., 1978). This forms the crossing zone. Central to the design of the roundabout merging zone is yielding behaviour required of all road users at each roundabout approach. For this, a yielding-specific PET, the yielding post-encroachment time (yPET), is devised by using the overlapping lanes of the merging zone as a predefined crossing zone. In this way, yPET remains comparable to any other PET involving yielding.

Much like in ordinary gap acceptance scenarios (Ahmed, 1999), where a road user wishes to cross a stream of traffic, a road user, upon yielding, accepts a certain gap between a lead and a lag road user before entering a traffic stream with the right of way (the conflicting flow). Thus, approaching road users at a merging zone are associated with two yPET measures: one in relation to the leading vehicle, the lead yPET; and one in relation to the lagging vehicle, the lag yPET. Note that these yPETs, or gap acceptance measures, are not to be confused

¹This is in fact the reference implementation included with tvaLib.

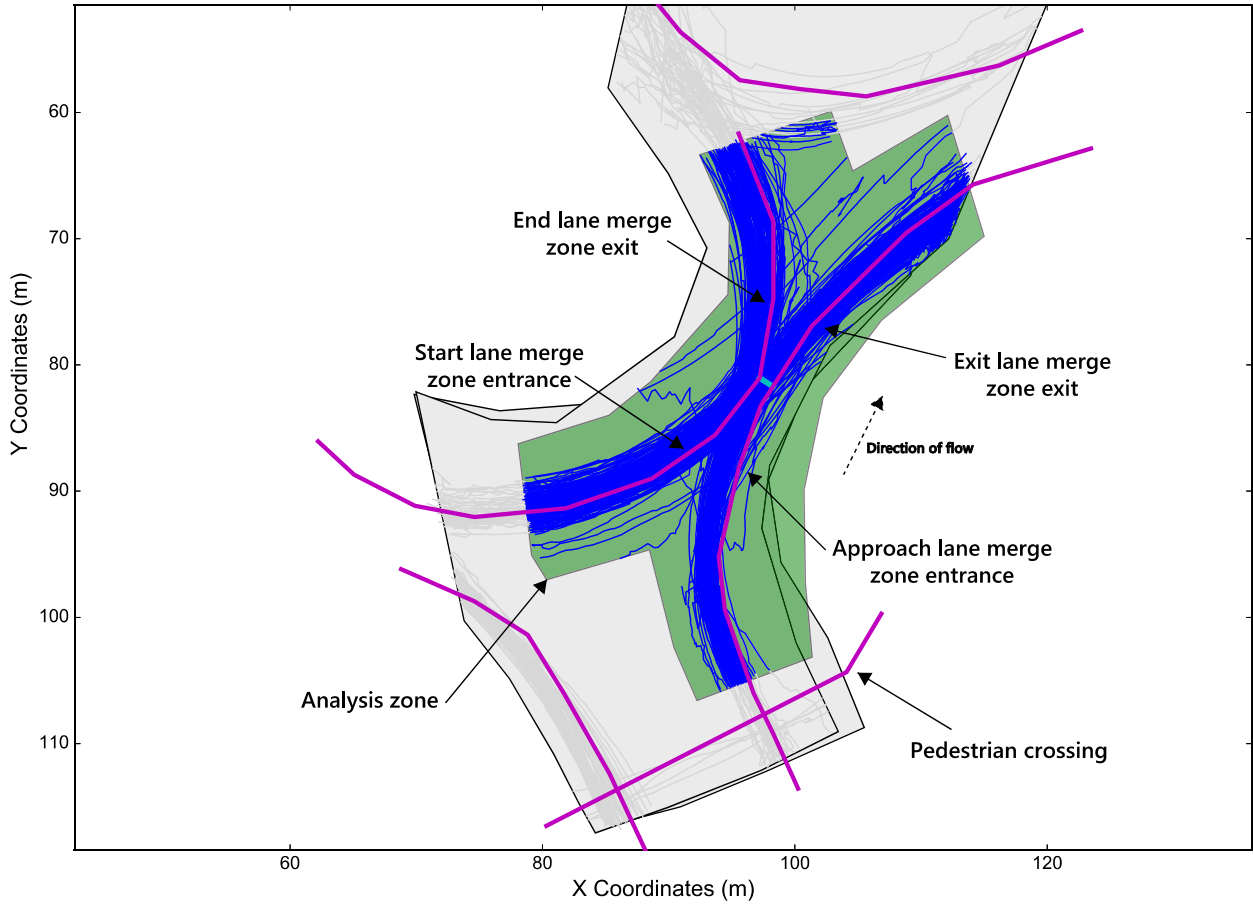


Figure 5.3 Roundabout merging zone metadata supplied to the HLI roundabout module.

with the GT as used in the literature, as this measure makes use of motion prediction.

These $yPET$ s are measured as the known, not predicted time difference between an approaching road user entering the merging zone from an approach lane at the instant t_{app} , and the next or previous conflicting flow road user entering the merging zone from a conflicting lane at times t_{prev} or t_{next} respectively. Thus

$$yPET_{lead} = t_{app} - t_{prev} \quad (5.1)$$

and

$$yPET_{lag} = t_{next} - t_{app} \quad (5.2)$$

This measurement is illustrated in Figure 5.4. The lead gap is measured when the conflicting flow road user leads the approach road user through the merging zone, while the lag gap is

measured when the conflicting flow road user follows the approaching road user through the merging zone. In theory, and in driving practice, motorists on the approach lane must yield to motorists approaching the merging zone from the conflicting lanes. These rules are not always respected, however. In the data, this should be seen in the form of very small lead yPET measurements, which would be indicative of aggressive tailgating. Meanwhile, very small lag yPET measurements would be symptomatic of lack of yielding. On the other hand, yPETs significantly longer than saturation headways (about 2 s (Shao et al., 2011; Shao and Liu, 2012), though likely slightly longer, to account for yielding deceleration and acceleration action in congested flows) are likely the result of random arrivals (especially in unsaturated flow), e.g. Poisson arrivals.

The yPET measures calculated in this manner are generated from ordinary random arrivals indiscriminately, whether any yielding truly occurs. Those interactions without yielding are of little interest from a road safety point of view, since they do not have the precondition of involving two or more road users interacting simultaneously. To solve this problem, a threshold value is chosen to distinguish between yPET measures and ordinary random arrivals. Compared with ordinary saturation headways of 2 s (Shao et al., 2011; Shao and Liu, 2012), a conservative value of 5 s is chosen. In any case, a site ranking based on mean yPET below 5 s is shown to be highly correlated with a site ranking based on mean yPET below 3 s (Spearman's correlation of 0.8029 based on a sample of 35 sites), thus suggesting that headway thresholds of 3 s or 5 s are interchangeable for yPET. The full correlation table is provided in Table 5.1.

Table 5.1 Spearman's Correlation of Ranking Based on Mean Site yPET below Threshold Value

	$\overline{yPET}_{<1s}$	$\overline{yPET}_{<3s}$	$\overline{yPET}_{<5s}$
$\overline{yPET}_{<3s}$	0.3607		
$\overline{yPET}_{<5s}$	0.4686	0.8029	
$\overline{yPET}_{<10s}$	0.2941	0.4897	0.8374

Movements

Given that a merging zone has two origins and two destinations, a total of four distinct movements can be expected from road users. These movements are, in order,

- (i) road users travelling through the merging zone on the conflicting lanes exclusively, i.e. within the roundabout ring only,

- (ii) road users leaving the roundabout conflicting lanes via an exit lane,
- (iii) road users entering the roundabout conflicting lanes via an approach lane, and
- (iv) road users entering and exiting the roundabout immediately, via an approach lane and the next exit lane.

These movements are presented in Figure 5.5. As these are movements defined at the merging zone scale, movement through the entire roundabout is normally a sequence of chained merging zone movements. For a typical four-way roundabout, a right turn at the roundabout involves movement (iv) only; a through movement corresponds to movement (iii) followed by movement (ii); a left turn (or a U-turn) corresponds to movement (iii) followed by one or more movement (i), and finally one movement (ii).

Regardless of the number of exits and where a road user wishes to exit the roundabout, the movement type across the merging zone can be generalized following these rules. Therefore, some roundabout HLI analysis is aggregated according to these movement types.

Flow Ratio

Table 5.2 lists the exposure factors calculated by means of HLI. These include the flow ratio Q_r , expressed as the ratio between the difference in approaching flow Q_{app} and conflicting flow Q_{conf} (i.e. flow in the start lane of the merging zone), and the sum of these flows:

$$Q_r = \frac{Q_{app} - Q_{conf}}{Q_{app} + Q_{conf}} \quad (5.3)$$

The flow ratio describes two phenomena:

- As Q_r tends towards 0, flows Q_{app} and Q_{conf} approach equal size, resulting in balanced merging demand. In a roundabout scenario, however, as the flow Q_{conf} has priority, if gaps are too small, even balanced flows may still result in merging issues if total demand is large. As Q_r tends towards 1 or -1, either Q_{app} or Q_{conf} respectively tends to dominate merging action. The tendency towards equal flows or dominance of any particular branch is expressed with the absolute flow ratio $|Q_r|$.
- The sign of Q_r indicates which branch dominates flow. When the approach flow Q_{app} dominates, and when normalized flow ratio Q'_r , defined as

$$Q'_r = \frac{Q_r + 1}{2} \quad (5.4)$$

and ranging from 0 to 1, approaches 1, then, assuming relatively low or moderate flows overall, yPET and TTC measures are expected to be high (and safe), as approaching road users should expect large and frequent gaps in the conflicting flow. However, in the reverse case, as Q_{conf} dominates and Q'_r approaches 0, then, assuming overall flows are significant, yPET and TTC measures are expected to decrease (and be less safe) as road users struggle to find gaps to merge into traffic.

Furthermore, the inflow Q per hour per lane is given as

$$Q = \frac{Q_{app}}{n_{app}} + \frac{Q_{conf}}{n_{conf}} \quad (5.5)$$

where n_{app} and n_{conf} is the number of approach and conflicting lanes serving traffic, respectively.

Table 5.2 Indicators Characterizing Demand and Traffic Ratios for HLI at Roundabouts

Variable (aggregated to the unit of analysis)	Factor description	Type (Units)
<code>flowratio</code>	Ratio Q_r of approaching and conflicting flows	Numerical $\in [-1, 1]$
<code>absflowratio</code>	Absolute value $ Q_r $ of the ratio of approaching and conflicting flows	Numerical $\in [0, 1]$
<code>approach_dominance</code>	Approach flow dominance Q'_r	Numerical $\in [0, 1]$
<code>approach_flow_ph</code>	Hourly approach flow Q_{app}	Numerical (veh/h)
<code>conflicting_flow_ph</code>	Hourly conflicting flow Q_{conf}	Numerical (veh/h)
<code>inflow_phpl</code>	Hourly traffic volume Q normalized for number of lanes	Numerical (veh/h)

5.2 Québec Roundabout Data Inventory

This chapter makes use of the roundabout inventory collected as part of the Québec Roundabout Safety research project (proposal number 2012-SO-163493) (Saunier et al., 2015; St-

Aubin et al., 2013a; Burns et al., 2013). This data comprises a near comprehensive inventory of all 145 roundabouts in the province at the time of data collection (2014). Data collection for that study was performed following identical measurement standards outlined in section 5.2.3 and are thus essentially purpose-built for this work. The relatively low number of roundabouts present in this province, despite a territory of 1,542,056 km² and a population of over 8 million inhabitants, can be explained by a concentration of urban density along the southern border with the United States and roundabout construction being a relatively new phenomenon. The first roundabout in Québec was built in 1998, but the majority of the 145 roundabouts have been built in the last 5 years. Nevertheless, Québec has the highest number of roundabouts in Canada.

5.2.1 Site Selection

Figure 5.6 maps the location of all known roundabouts across the province of Québec as of 2014, as well as those roundabouts targeted for safety analysis with video data collection. These roundabouts are selected using a ten-point scale with weighted average scores for multiple factors, including regional distribution, representation of various administrative bodies (who implement and manage roundabouts differently, i.e. the MTQ, responsible for highways and regional inter-city rural highways, and municipalities), spatial clustering (to reduce data collection costs), traffic density,² and feasibility of performing video data collection (according to specifications laid out in section 3.2.1 and with permission from the local authority). See appendix D for details of this process. Ultimately, 23 roundabouts were visited over a period of 3 years (2012-2014). At each site, a sample of one day of typical traffic is recorded, starting anywhere between 6:00 a.m. and 7:00 a.m., and ending anywhere between 6:00 p.m. and 10:00 p.m.

In order to control for visibility conditions and seasonal, weather, and societal effects, data collection timing is carefully planned. Most video recording is performed under partially overcast or moderately sunny weather conditions (though recording at one site did witness moderate rain) during daylight hours. This has the added benefit of controlling the tracking quality; as mentioned in Chapter 3, video analysis is sensitive to the quality of lighting conditions. Furthermore, no significant sunlight blinding is included in the data collected at the sites visited. Meanwhile, care is taken so as to only record video data during statutory

²Very low traffic volumes produce too few observations. There must be at least two road users inside the analysis zone simultaneously to produce a single interaction observation. An average rate of ten road users per minute generates single-digit user pairs per hour, only a fraction of which will generate a meaningful SSM.

work days during no significant vacation periods, in no proximity of a construction zone. Road surface conditions are always clear of snow, though it rained lightly on some days, and heavily only for very short periods.

At each site, one to three VCUs are deployed, depending on availability of equipment, to cover merging zones as far as possible. A total of 52 camera views (a number of video sequences recorded from the same stationary camera, in sequence) are recorded, each covering one or more merging zones. A few of these camera views are rejected, while some others have overlapping views on the same merging zone. Ultimately, this process leaves 35 analysis zones.




Figure 5.4 Time instants for calculation of yielding post-encroachment time for a roundabout merging zone.




Figure 5.5 Roundabout merging zone movement types.

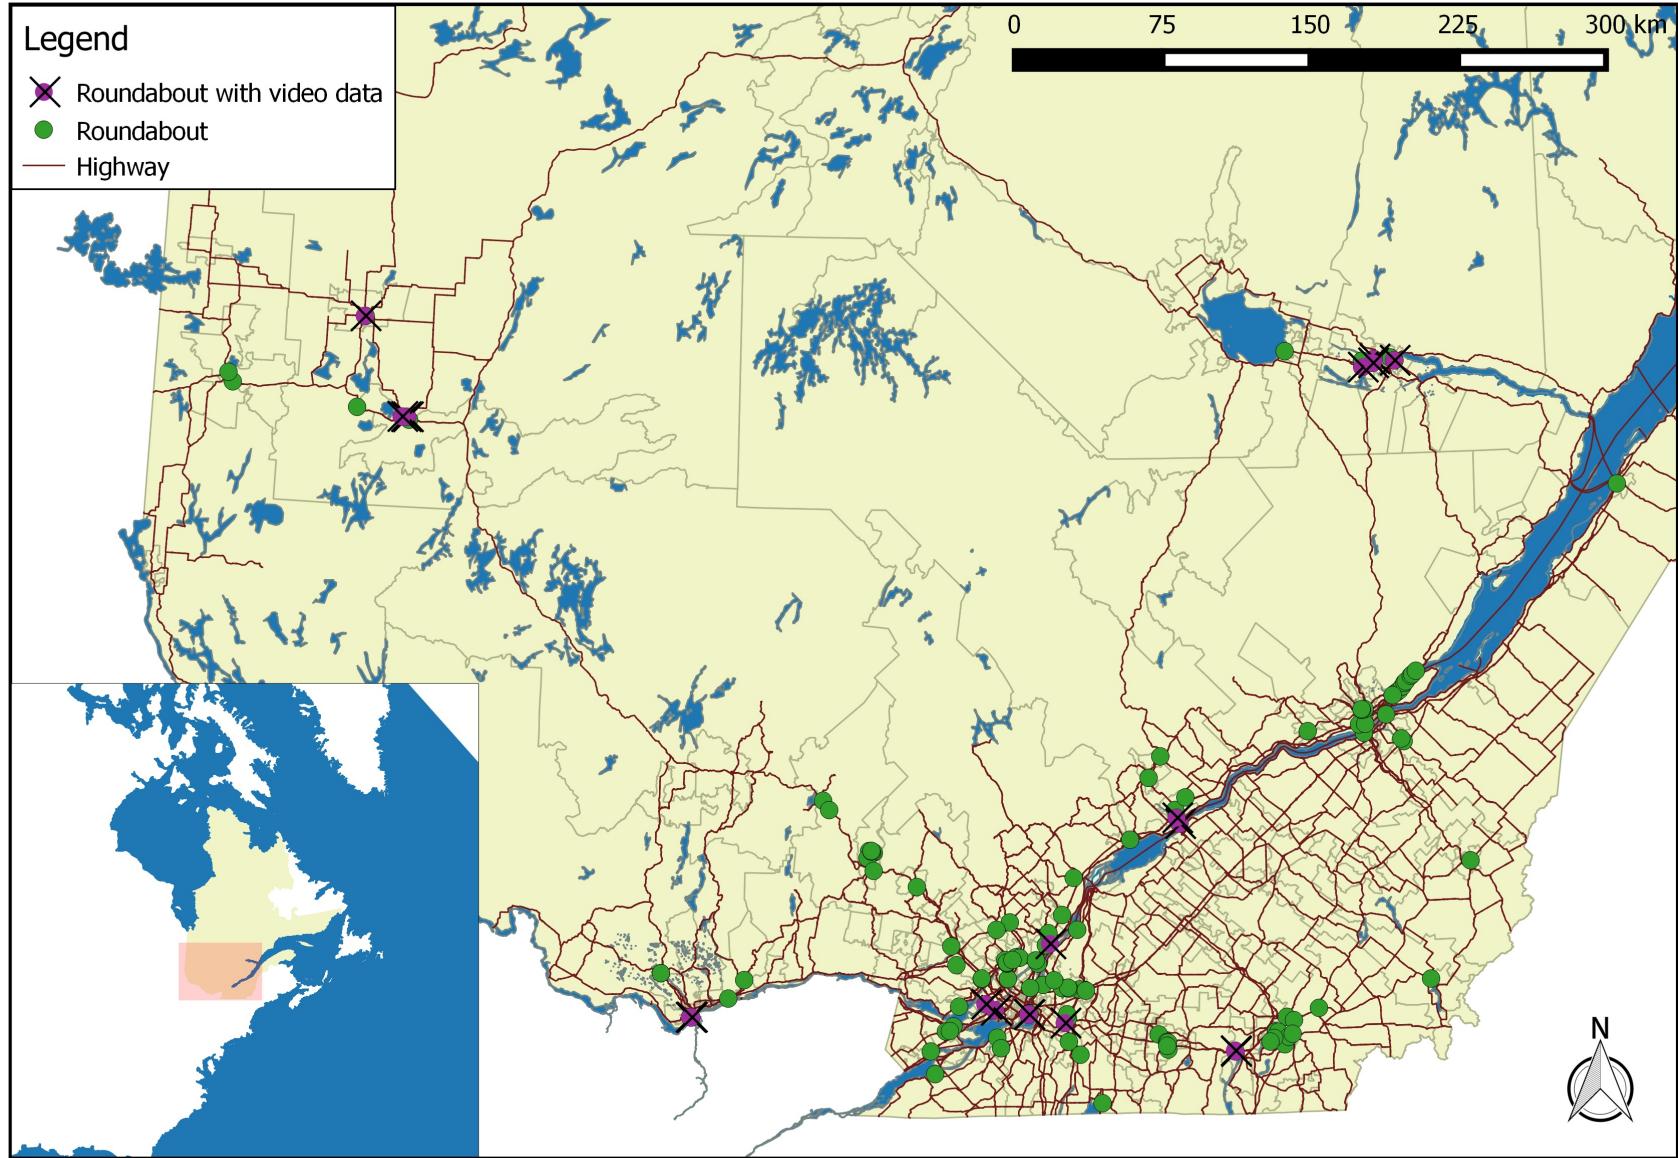


Figure 5.6 Roundabout locations across the province of Québec.

5.2.2 Data Inventory

A summary of the size of the data can be found in Table 5.3. The sample size is 196,808 unique road users. The vast majority are passenger vehicles, and the remainder is composed largely of buses and trucks. Given that there are 5,241,846 registered drivers in the province of Québec (SAAQ, 2014) and that data collection took place over the course of no more than a single day at each site, and assuming that each driver appears on average exactly once,³ then it can be concluded that the sample is representative of approximately 3.7% of the entire Québec population. This subjects the conclusions to the limitation that the road users captured in the study live and commute near the roundabout and are probably much more familiar with it than other road users. On the other hand, this is a useful property when comparing road users between regions, as in Chapter 6).

Table 5.3 Québec Roundabout Data Inventory

Data collection		
Roundabouts visited	23	
Camera views	52	
Merge zones/analysis zones	35	
Total hours of video (pre-analysis)	534.2 h	
Disk space usage		
Video data	1,518,854 MB	(2,568 files)
Trajectory data	289,411 MB	(909 files)
Cached data	114,516 MB	(5,089 files)
Misc. overhead (e.g. system)	21,858 MB	(12,596 files)
Trajectory Data		
Unique road users	196,808	
Vehicle-kilometres travelled	11,519.1 veh-km	
Duration (analysis)	435.6 h	
Interaction Data		
User pairs	176,749	
User pairs with TTC (cmp, 15t)	32,650	


5.2.3 Potential Contributing Factors

Diagrams of the basic potentially significant contributing factors of road safety of roundabout merging zones which were collected are presented in Figure 5.7. These include predomi-

³Although it is reasonable to conclude that a driver would take the same route twice in a day as part of commuting, each merging zone in which the driver appears covers only a single direction of travel through a roundabout.

nantly basic roundabout design features (Figure 5.7a) and geometric design measurements (Figure 5.7b). As part of the inventory prepared for the Québec Roundabout Safety research project (Saunier et al., 2015; St-Aubin et al., 2013a; Burns et al., 2013), additional factors are measured, including signage, pedestrian features, approach angles, and curb lengths (for a total of 60 potential contributing factors), but ultimately the sample of roundabouts with video data is not large enough to investigate so many factors. In any case, few pedestrian observations had been made, so factors related to pedestrians are not included in the final study presented here. However, pedestrian safety has been studied in similar, related work (Perdomo et al., 2014; Perdomo, 2015).

Other geometric factors omitted for this study, but which may be of interest in a future study, include approach trajectory deflection, centre island height (Jensen, 2014), and general measures of approach visibility. Centre island height had not been considered until after data collection took place. Trajectory deflection had been considered, however proved very difficult to define and parameterize (beyond the practice of using inside radius and apron width) as roundabout approaches sampled in Québec were rarely straight, and even less consistent in design.



(a) Basic design features

(b) Basic geometric measures

Figure 5.7 Basic geometric design elements of a roundabout merging zone.

Table 5.4 summarizes the geometric and environmental factors collected for this study. These potentially contributing factors exist at either the level of a merging zone or the level of a roundabout.

The factors collected at the merging zone level are almost all numerical measurements of geometric features. These measurements were taken either on-site or through mostly recent, high-resolution aerial photography. These features mainly encompass aspects of lane configuration, roundabout radii, and speed limits (ranging between 30 and 90 km/h, with all legally sanctioned speed limit intervals represented), but they also include curb typology (painted chevrons versus poured concrete), the presence of private driveways, and the upstream distance from the nearest intersection to the approach (which has been shown to have an effect on speed as demonstrated in Hydén and Várhelyi, 2000).

The factors measured at the roundabout level are all environmental and qualitative in nature and collected manually (on-site visit). These encompass the type of land use, the density of the built environment, and a basic classification of the type of network that the roundabout is connected to. In almost all instances, the roundabout environment classification is homogeneous on all sides of the roundabout. In the few cases where it is not, the predominant environment is used. Land use is classified as follows:

- Land use type 1 (**lu1**) is characterized by **vacant land** devoid of any use other than transportation. This is typical of roundabouts operated by the MTQ, as they oversee regional inter-city transportation facilities on government-owned land in rural areas, e.g. highway facilities and support facilities such as access ramps and interchanges.
- Land use type 2 (**lu2**) is characterized by **residential** land use of all types, including detached housing and multi-unit housing.
- Land use type 3 (**lu3**) is characterized by **commercial** land use of all types, including big-box stores, office space, and small, dense commercial strips.
- Land use type 4 (**lu4**) is characterized by **industrial** land use of all types, from manufacturing to resource extraction.
- Land use type 5 (**lu5**) is characterized by **mixed** land use. In practice, this typically means a roughly equal mix of residential and commercial land use. It should be noted that the exact proportion could vary somewhat. Most sites characterised by this land use are situated along commercial arterials within dense residential neighbourhoods. Mixed land use is common in these highly walkable, urban neighbourhoods.

- Land use type 6 (**lu6**) is uncommon. It is **institutional** land use, including a range of public services such as hospitals, schools, and other government-agencies.

Roundabouts surrounding road networks are classified according to the designation of the primary traffic corridor passing through it (with the highest AADT). Roundabouts can be found attached to the following types of road networks:

- Network class type 1 (**nc1**) is the smallest and most basic road network: the roundabout is part of a road network with only **collector roads**.
- Roundabouts with a network class type 2 (**nc2**) designation can be found attached to at least one arterial road, including avenues and boulevards.
- Roundabouts with a network class type 3 (**nc3**) designation can be found attached to at least one regional, undivided highway (e.g. with speed limits of 70 to 90 kmh).
- By design of a limited-access-highway, roundabouts cannot be attached directly to a limited-access-highway. However, they can be found near limited-access-highways serving as interchanges with on-ramps and off-ramps feeding into them. These roundabouts are designated network class type 4 (**nc4**).

For the purposes of this study, urban density is measured very roughly from the buildings in the vicinity (about 1 km) of the roundabout. Thus urban density can be classified into the following four groups:

- Roundabouts with a density type 1 (**d1**) designation have no buildings at all in the vicinity. There is significant overlap of this type of roundabout with roundabouts with land use type 1 and type 4 designations (mostly composed of rural roads and inter-city roads, respectively).
- Density type 2 (**d2**) is characterized by detached housing; small, single-story businesses; or farms.
- Density type 3 (**d3**) is characterized by semi-detached housing; medium-sized businesses; heavy-industry; or institutional land use.
- Density type 4 (**d4**) is characterized by multi-story buildings (10 or more). Given the large footprints of roundabouts, these are exceedingly rare in very dense environments.

Table 5.4 Geometric Factors

Variable (merge zone)	Factor description	Type (Units)
b_quad_type	Merge zone only contain an exit ⁴	Categorical
n_start_lanes	Number of start (conflicting) lanes	Numerical
n_end_lanes	Number of end (conflicting) lanes	Numerical
n_app_lanes	Number of approach lanes	Numerical
n_exit_lanes	Number of exit lanes	Numerical
n_slip_lane	Number of slip lanes	Numerical
b_app_med_type	Median type (0=raised, 1=painted)	Categorical
b_driveway	Presence of a driveway on merging zone	Categorical
a_quad_size	Angular size of merge zone	Numerical (°)
r_out_start	Outside diameter at start of merge zone	Numerical (m)
r_in_start	Inside diameter at start of merge zone	Numerical (m)
r_out_end	Outside diameter at end of merge zone	Numerical (m)
r_in_end	Inside diameter at end of merge zone	Numerical (m)
w_apron	Width of apron	Numerical (m)
w_lane1	Width of lane 1	Numerical (m)
d_app_inter	Upstream distance to nearest intersection	Numerical (m)
app_speed_limit	Mandatory speed limit on approach	Numerical (km/h)
Variable (roundabout)	Factor description	Type (Units)
land_use	Land use classification	Categorical
network_class	Classification of surrounding road network	Categorical
density	Urban density classification	Categorical

5.3 Overview and Exploratory Analysis

In the following sections, descriptive statistics are presented for traffic flows, speed, yPET and TTC measures across the 35 merging zones being studied.⁵

For the most part, individual traffic (road user) measures of speed and yPET at each site follow normal distributions for that site and are aggregated at the analysis zone/merging zone level using an average statistic. TTC measures on the other hand are not quite normally distributed and, furthermore, incorporate road user level factors; they are thus examined at a disaggregated level.

⁵These statistics are automatically generated through tvaLib or tvaLib's included roundabout HLI module.

5.3.1 Flows

The hourly approach flows (`approach_flow_ph`) range from 19 vehicles/h to 1,590 vehicles/h, with mean and median values of 453 and 423 vehicles/h respectively. Hourly flow aggregated to the site is not normally distributed, but 80% of the hour-long sequences studied have between 75 and 950 vehicles/h. The number of hourly conflicting flows (`conflicting_flow_ph`) shares a similar pattern, though the minimum number of observations is 41 vehicles/h and the mean and median values are a bit lower, at 373 and 265 vehicles/h respectively.

Twenty-five of the 35 sites have a flow ratio slightly favouring the approach, with a mean and median `approach_dominance` of 0.547 and 0.509 respectively (0.5 is perfectly balanced). The mean and median `absflowratios` are 0.372 and 0.353 respectively, suggesting that flow ratios between the approach and the exit are more balanced (at 0) than unbalanced (at 1). In other words, approach demand tends more towards proportionality with conflicting flows than not, but not perfectly. In summary, there is reasonable flow ratio representation present in the data, with minimum and maximum `approach_dominance` values of 0.147 and 0.920 respectively.

Figure 5.8 plots the sequential arrivals of road users at the approach, and at the start of the merging zone, across all sites. Approach sequential arrivals, i.e. platoons, are arrivals which are performed in succession without interruption from conflicting flows, and vice versa for conflicting flow sequential arrivals. There is no maximum time criterion. Overall, regardless of origin, sequential arrivals of road users of more than two or three at a time are very uncommon, again suggesting balanced arrivals. It is somewhat surprising to see, overall, a greater frequency of single arrivals (at a time) occurring for the conflicting flow, even though it has priority in the roundabout. This effect can be seen at individual sites with the heaviest flows, however, as in Figure 5.9a. It should be noted that a bit more variety of distribution exists when individual sites are observed, though overall, a similar trend can be found.

When maps of trajectory positions are examined, an interesting pattern emerges in path clusters taken by road users. At a large number of single-lane roundabouts, each distinct movement, as defined in section 5.1.2, occupies a distinct area of the merging zone, not unlike the exaggerated paths outlined in Figure 5.5. Movements originating from a conflicting lane and heading towards an exit (movement type (ii)), or movement originating from an approach and heading towards a conflicting lane (movement type (iii)), form a crossing pattern within the merging zone. Movements within the ring or entering and immediately exiting, types (i) and (iv) respectively, form side-by-side paths on the inner side and outer side of the merging zone, respectively. However, this pattern, illustrated in Figure 5.10a, can only be found in single-lane roundabout merging zones of small angular size (`a_quad_size`) and where lane




Figure 5.8 Sequential arrivals of road users at the approach, and start of all merging zones.

widths (w_{lane1}) are sufficiently large.

A counter-example where this effect cannot be witnessed is illustrated in Figure 5.10b. Thus at some sites, good or perfect sequential mixing (full lane changes into the conflicting flow) of road users is taking place, whereas at others, only reduced mixing is taking place. The implication of this observation is that road users are moving side-by-side across these merging zones where reduced mixing is occurring, despite a single-lane design intent. And indeed, upon manual inspection of these trajectories, it appears that road users are treating large lane widths as *de facto* slip lanes.

Whether or not this has significant implications for any aspect of road safety—one way or another—is not entirely clear. It is difficult to parametrize these trajectory patterns for statistical analysis, particularly because they occur to different degrees, however the effect appears to be correlated with, not surprisingly, merging zone angular size and, to a lesser degree, lane widths, two factors both already parametrized for regression models in section 5.4. In any case, this analysis illustrates the benefits of surrogate safety methods

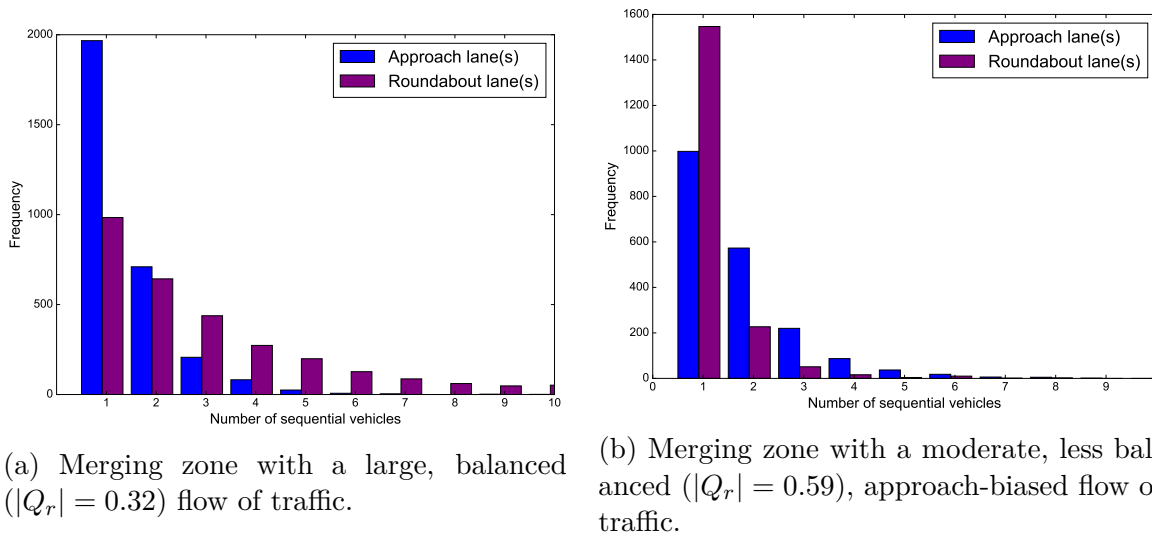


Figure 5.9 Sequential arrivals of road users at the approach and start of individual merging zones.

as not only a potential substitute for quantitative historical accident analysis, but also as a better approach to understanding traffic interactions at the microscopic level, and ensuing collision mechanisms.

5.3.2 Speed

Across individual lanes, mean road user speeds are almost universally normally distributed, except on some approach lanes, as some road users slow or stop to yield, while others do not. This is exemplified in the average speed data plotted by lane in Figure 5.11 (where each lane is represented by one alignment). This speed data represents the average of each road user's speed measured in a corresponding lane, itself an average of the individual speed measurements made at that road user's positions associated with the corresponding alignment. When merging speed is measured from different lanes, results tend, not surprisingly, to resemble a mixture of normal distributions. The aggregate of all speed values across all sites has a normal distribution.

Figure 5.12 plots the spatial distribution of average speed at two typical roundabout merging zones. Figure 5.12a depicts a typical merging zone with a relatively balanced flow ratio ($|Q_r| = 0.32$, favouring the conflicting flows slightly). Characteristic features are visible: reduced speeds are present on the approach (5–10) km/h, just before the merging zone, as part of yielding action by drivers in this lane; road users accelerate to 40–45 km/h upon exiting the merging zone and roundabout; and road users operate at a constant speed, ranging between

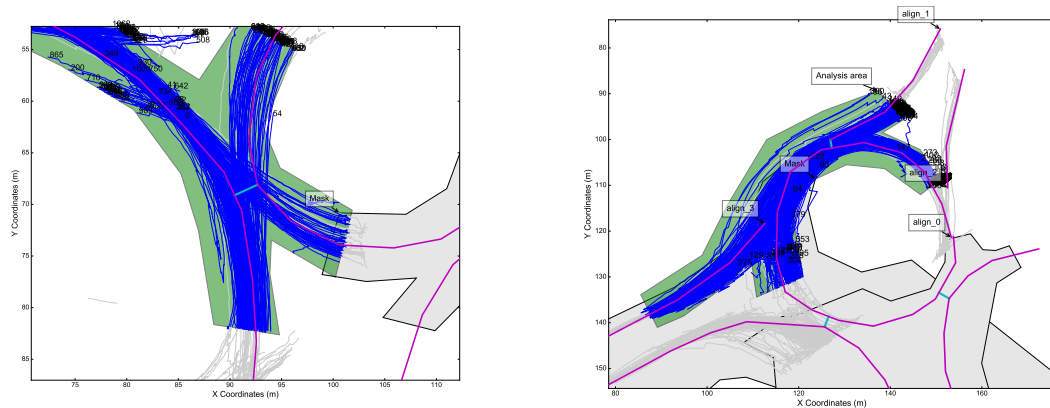
(a) Merging zone with angular size of 90° (b) Merging zone with angular size of 145°

Figure 5.10 Road user trajectories (blue lines) across the merging zone (green area) for two sites of differing angular size.

25 and 35 km/h, within the roundabout and merging zone. Two important details are of note here: firstly, the inside edge of the roundabout merging zone, next to the central island, hosts slightly higher speeds than the outside edge, which may be symptomatic of reduced deflection action, and secondly, there is a mismatch in speed of about 5 km/h between movement types at the exact point of crossing in the merging zone, especially visible in Figure 5.12a. This is in contrast to the roundabout merging zone depicted in Figure 5.12b, which has an unbalanced flow ratio heavily favouring the approach. Speeds across this merging zone are much higher, with little sign of slowing, as traffic volumes tend to be low, and more importantly, there are no conflicting flows to yield to.

Speed distributions (in the format of box plots) at the instant road users enter the merging zone are provided by site and by movement type in Figure 5.13. Note that sites lacking a box plot have too few road user observations of that particular movement type to be included (minimum of at least 50 road user per movement type). Several things can be concluded from this data:

- Movement types (iii) and (iv), which both involve approaching road users, have median and average speeds in the neighbourhood of 30 km/h, and rarely outside of the range of 20–40 km/hh.
- There is no correlation between number of observations and movement type, suggesting that a large variety of merging zone usage cases are represented here.




Figure 5.11 Speed distributions by lane (alignment), at a typical multi-lane merging zone with $|Q_r| = 0.59$.

- Median speeds for movement types (iii) and (iv) are roughly correlated with each other, but not nearly as much with and between movement types (i) and (ii).
- Between individual sites, a large degree of speed variability exists, both between quartiles and extreme values (a handful of these roundabouts are in 70 or 90 km/h zones). Note that outlier observations are not plotted in Figure 5.13, for the sake of clarity.

The speed measures provided in Figure 5.13 are taken at the instant when a road user enters the merging zone.

Entire speed profiles have been investigated throughout the entire length of the merging zone. The general trend is that speeds tend to increase, not surprisingly, as road users leave the merging zone, and especially the roundabout entirely, i.e. using the exit lane. This is illustrated in Figure 5.14, where relative distance is normalized with respect to the start and end of the merging zone. Each speed value is aggregated to a normalized segment; there are 20 segments within each merging zone. Aggregated, there is an average difference of up to 10 km/h between movement types across significant portions of the merging zone and beyond (St-Aubin et al., 2013a).

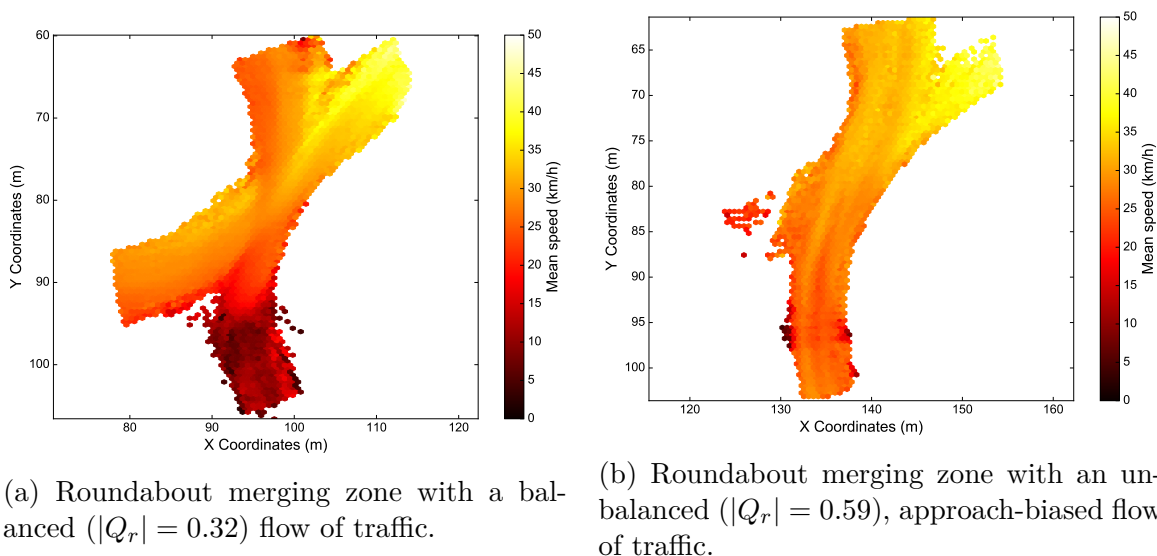


Figure 5.12 Spatial distribution of average speed for roundabout merging zones throughout an entire day. In both cases, road users approach the merging zone from the south-west. All speeds are normalized between 0–50 km/h.

5.3.3 Yielding Post-Encroachment Time

yPET measures are calculated for all sites. The aggregated frequency of all lead and lag yPETs across all sites is presented in Figure 5.15 for values up to 5 s. Both distributions appear to be composed of a mixture of a normal distribution and a negative exponential distribution (which is traditional for headways and hence gaps), plus a spike of values near $yPET = 0$ which may be oversampled with tracking errors.

Two things are of particular note. First and foremost, most yPET measures where the approach road user follows a conflicting road user (approach road user lead yPET) are clustered around a peak of 2.2 s, roughly between 1 and 3 s. This is just slightly larger than the car following saturation headway distribution mean and median values found in the literature (Shao et al., 2011; Shao and Liu, 2012), plus an additional mean time of about 0.2 s, which may be explained by lost time due to deceleration and acceleration at the approach. Approach road user lag yPET also peaks close by, at a slightly lower value of 1.8 s. Lag road users, constituting the conflicting flow, are not required to yield. This suggests that, for the most part, proper yielding is occurring, with normal gaps during merging.

As for the decreasing yPET measures beyond the values for saturation headway, this is to be somewhat expected, since some of the yPETs are purely the result of random arrivals between individual road users, though the lowest values are of some concern. yPET observations below




Figure 5.13 Speed quartiles for each site by movement types (i) through (iv). Speed includes observations during all times of the day. Sites are ordered by the median speed of movement type (iv). Numbers at the bottom of each graph are the sample size for each site (minimum 50 to be plotted).

1.5 s are of special concern for safety, particularly regarding SEC methodology.

However, on a site-by-site basis, lead and lag yPET distributions appear more dissimilar to each other; this may be explained by site-specific factors. Furthermore, a large disparity between gap distributions can be found between sites. The box plots of lead yPET measures for each site in Figure 5.16 show this disparity. The box plots of lag yPET measures for each site show a near identical trend, though the order of sites by median yPET is different. The correlation of median lead yPET below 5 s and median lag yPET below 5 s is found to be only 0.1687.




Figure 5.14 Aggregated mean speed and standard deviation profiles by movement types (i) through (iv).

5.3.4 Time-to-Collision

As discussed in section 4.6, TTC safety indicators based on discretized motion pattern prediction and 15th centile aggregation (TTC_{15}) are chosen for evaluation, due to their intrinsic benefits in non-linear driving environments, as is the case with roundabouts. From Table 5.3, it can be shown that 18.5% of user pairs (8.3% of road users) are involved in at least one collision course at a roundabout merging zone. Furthermore, the total expected wSE (for SEC methodology) is found to be 2,168, or 1.2% of user pairs (1.1% of road users). Recall that wSE is probabilistic in nature: this is the number of *expected* serious events.

An exhaustive fit of 82 different probability distributions is performed on the sample round-




Figure 5.15 Yielding post-encroachment time distributions, pooled across all sites.

about data using Maximum Likelihood Estimation as implemented in SciPy (Jones et al., 2001–). This is reduced to a list of 41 probability distribution models with very similar fit, all having residuals between estimated and observed values $< 5.5\%$ different from one another (see Figure D.1). Two families of distributions emerge: normal type and gamma type. Since most of these models are essentially identical, only the popular gamma, normal, loggamma, and lognormal are kept. Figure 5.17 plots the probability distribution of TTC_{15} safety indicators and candidate probability distribution functions fitted to the data. Of the four remaining models, the lognormal model,

$$P(TTC_{15}) = \frac{1}{TTC_{15}\sigma\sqrt{2\pi}} e^{-\frac{(\ln TTC_{15} - \mu)^2}{2\sigma^2}} \quad (5.6)$$

with location parameter $\mu = 0.302$ and scale parameter $\sigma = -1.731$, and the gamma model,




Figure 5.16 Lead yielding post-encroachment time box plots for each site, ordered by median.

$$P(TTC_{15}) = \frac{1}{\Gamma k \theta^k} TTC_{15}^{k-1} e^{-\frac{TTC_{15}}{\theta}} + e \tag{5.7}$$

with parameters shape $k = 4.560$, scale $\theta = 0.638$, and offset $e = -0.245$, have the best Kolmogorov–Smirnov statistics at 0.058 and 0.053 respectively, although they each have slightly higher residuals than loggamma and normal distributions. For reference, the fitted normal distribution has a mean $\mu = 2.663$ and a standard deviation $\sigma = 1.321$. Another issue with the normal distributions is that TTC values below 0 are impossible. This normal distribution fitted to the data predicts about 5% of TTC_{15} safety indicators with a negative value.

Normal distribution fitting is more forgiving at individual sites, as demonstrated in the curve fitting for the three sites with most user pairs, illustrated in Figure 5.18. It should be noted that Figures 5.18a and 5.18c appear to show the possibility of bimodal and trimodal distributions respectively. However, it is important to recognize that sample sizes at individual sites tend to be smaller, so it is difficult to draw conclusions about distribution shapes from




Figure 5.17 Curve fitting for probability distribution of all TTC_{15} values.

individual sites. In any case, it might be interesting as part of future work to examine mixture distributions at individual sites, perhaps by separating TTC observations by interaction type, e.g. by angle of incidence.

Angles of Incidence

The distribution of median angle of incidence for each user pair is presented in Figure 5.19. This contrasts with the distribution of angles of incidence of only those user pairs with a corresponding TTC measure. Overall, the general shape of distributions is as expected for a roundabout: low angles suggest most interactions are of the rear-end conflict type or, to a lesser degree, of the -variety, given that all merging zone interactions result from car-following or from converging manoeuvres.

However, it is interesting to note the discrepancy in shape of the distributions of angles of




Figure 5.18 Curve fitting for probability distributions of TTC_{15} values at top three sites with most user pairs.

incidence, the smaller ones in particular. While the distribution of median angle of incidence generally follows a lognormal-like distribution, the distribution of median angles of incidence corresponding to only those user pairs with a measured TTC value shows an increased concentration of rear-end conflict-generating or side-by-side-side-swipe conflict-generating angles of incidence. This might be explained by the fact that a number of side-swipe conflict interactions are benign, or that, given the curved nature of merging zones, many of the interaction instants occur at different angular locations within the merging zone.

Cluster Analysis

Due to the large number of factors, an initial cluster analysis of sites by their geometric factors is performed for the purpose of road safety analysis (St-Aubin et al., 2015c). K-means clustering is employed on all of the variables. Several trials are performed for between three and six centroids to find a suitable segmentation that (i) produced meaningful and interpretable clusters, and (ii) produced a regression model with explanatory power and




Figure 5.19 Distribution of median angle of incidence per user pair. The thick line plots the distribution for user pairs with at least one TTC observation.

statistical significance. These clusters are used in the statistical model of road safety and road user behaviour to follow, using the most up-to-date SSM calculations. Table 5.5 lists these clusters, provides a short description for each, and provides sample size statistics (group size and number of user pair observations with safety indicators in each).

From Figure 5.20, it can be seen that clusters cl_4 (residential collectors with low absolute flow ratios and short upstream distances to nearest intersection) and cl_6 (two-lane arterials near commercial or institutional land use and very high approach flow dominance) offer clear and unambiguous shifts in TTC distribution, and hence theoretical reductions to collision probability as per the SCC analysis methodology proposed in section 4.6.1. The remaining clusters have inconclusive TTC distribution shifts.

Table 5.5 K-means Cluster Profiles

Cluster	Description	Group Size	Individual User Pairs
cl_1	Arterial with wide lanes, far distance to upstream intersections, and very low approach flow dominance, mixed land use	6	1388
cl_2	Regional, single lane highways in industrial complex with low absolute flow ratios	6	18609
cl_3	Mix of highway ramps and arterials with very high absolute flow ratios	13	3498
cl_4	Residential collectors with low absolute flow ratios and short upstream distances to nearest intersection	6	6370
cl_5	Traffic circle converted to roundabout (two lanes, extremely large diameters, tangential approach angle)	4	2056
cl_6	Two-lane arterials near commercial or institutional land use and very high approach flow dominance.	6	729

5.3.5 Analysis of Correlation of Factors

One of the objectives of the study is to explain aspects of road safety using specific geometry and land use parameters (exogenous variables), and not necessarily to predict the occurrence of traffic parameters or SSMs at the given sites (endogenous variables). Thus, any exogenous variable that can be explained by another variable included in the same model must be rejected. To this end, a correlation analysis is performed on the traffic, geometric, and environmental factors of Tables 5.4 and 5.2 to purge collinear factors for regression analysis of specific parameters in section 5.4. Of the geometric factors, r_{in_start} , r_{out_end} , and r_{in_end} are removed, as they are ultimately very highly correlated (> 0.95) with r_{out_start} . While some dissimilarities still exist, as it turns out, the variation in these measures taken identically at either side of the merging zone is not significant. Furthermore, it is known from the outset that

$$r_{out} = r_{in} + w_{apron} + w_{lane1}n_{start\ lanes} \quad (5.8)$$

w_{lane1} is found to be correlated with n_{start_lanes} , but not greatly so (-0.62). This reflects in part what appears to be a tendency that some planners have of building a fixed-width area of pavement for each merging zone and approach/exit element, regardless of the




Figure 5.20 15th centile time-to-collision pooled to each of the six merging zone clusters.

number of planned lanes.

Tables 5.6 and 5.7 show the correlation matrix for the remaining parameters.

Table 5.6 Correlation table for the remaining traffic, geometric, and environmental factors.

	app_speed_limitt	b_quad_type	n_start_lanes	n_end_lanes	n_app_lanes	n_exit_lanes	n_slip_lane	b_app_med_type	b_driveway	a_quad_size	r_out_start	w_appron	w_lane1	d_app_inter
app_speed_limitt	1													
b_quad_type	0.1119	1												
n_start_lanes	-0.5079	-0.1161	1											
n_end_lanes	-0.4699	-0.1161	0.7348	1										
n_app_lanes	-0.3848	-0.4361	0.5714	0.6987	1									
n_exit_lanes	-0.4427	-0.0934	0.6575	0.3643	0.4466	1								
n_slip_lane	-0.0998	-0.0294	0.2533	0.2533	-0.0811	-0.0934	1							
b_app_med_type	0.0087	-0.0422	0.0985	0.0985	0.393	0.1591	-0.0422	1						
b_driveway	-0.1782	-0.0525	0.0126	0.0126	0.0664	-0.1667	-0.0525	-0.0754	1					
a_quad_size	0.0847	-0.2364	-0.2883	-0.2682	-0.0105	0.0006	-0.2444	-0.0069	-0.0086	1				
r_out_start	-0.0892	0.0401	0.5442	0.5675	0.3695	0.2671	-0.0803	-0.1618	0.0386	-0.4152	1			
w_appron	0.6343	0.007	-0.3408	-0.3187	-0.0868	-0.1895	-0.0751	0.246	-0.183	0.1812	-0.3212	1		
w_lane1	0.424	0.2343	-0.6179	-0.3998	-0.417	-0.3012	-0.1934	-0.2453	-0.2381	0.1957	-0.0125	0.1718	1	
d_app_inter	0.2331	-0.0098	-0.276	-0.0569	-0.0134	-0.2245	-0.1177	0.0155	-0.1509	0.1834	-0.149	-0.0004	0.147	1
flowratio	-0.1004	0.2549	0.0377	0.1092	0.1139	0.2678	-0.0294	0.3756	-0.1576	-0.1355	0.1127	-0.147	0.1038	0.214
absflowratio	-0.2696	0.2497	0.051	0.0049	0.0238	0.377	-0.233	0.3749	-0.0534	-0.1219	-0.1453	-0.1397	-0.0627	-0.0597
approach_dominance	-0.1004	0.2549	0.0377	0.1092	0.1139	0.2678	-0.0294	0.3756	-0.1576	-0.1355	0.1127	-0.147	0.1038	0.214
approach_flow_ph	0.3041	0.4482	-0.0291	0.0308	0.0155	0.0898	-0.0384	0.5338	-0.2775	-0.3064	0.0777	0.4059	0.245	0.1001
conflicting_flow_ph	0.3591	-0.0915	-0.0604	-0.0061	0.0147	-0.1154	-0.0037	-0.1569	-0.2114	-0.0256	0.0264	0.5817	0.25	-0.0758
nc2	-0.6465	-0.1485	0.7817	0.6574	0.5462	0.491	0.198	0.0355	-0.0589	-0.2581	0.4242	-0.248	-0.4991	-0.2271
nc3	0.7751	-0.0934	-0.3685	-0.2219	-0.1167	-0.2963	-0.0934	0.1591	-0.1667	0.1274	-0.1213	0.6987	0.33	0.2192
nc4	0.2968	0.377	-0.3079	-0.3079	-0.3721	-0.2476	-0.078	-0.112	-0.1393	0.0685	0.0041	-0.0903	0.3826	0.1455
lu2	-0.4575	-0.1239	0.0296	0.0296	0.032	0.0369	0.2374	0.3408	0.2089	0.0723	-0.4911	-0.2803	-0.293	-0.0693
lu3	-0.0998	-0.0294	0.2533	0.2533	0.2738	0.3151	-0.0294	-0.0422	-0.0525	-0.0048	0.0494	0.2124	-0.2834	-0.0578
lu4	0.5907	-0.0857	-0.1846	-0.1846	-0.0887	-0.1021	-0.0857	-0.1231	-0.1531	0.0858	-0.0707	0.6023	0.2863	-0.1512
lu5	-0.2376	-0.07	0.4271	0.4271	0.3138	0.1667	-0.07	-0.1005	0.1667	-0.3916	0.8058	-0.4526	-0.1495	-0.2254
lu6	-0.3042	-0.0525	0.0126	0.0126	0.0664	0.0764	-0.0525	-0.0754	-0.0938	0.0865	-0.0937	0.1348	-0.0907	-0.0135
d2	0.0051	-0.198	-0.0355	-0.1599	0.0512	0.0589	0.1485	0.2132	0.2652	0.1022	-0.2433	0.1442	-0.092	-0.3414
d3	-0.335	-0.0857	0.4308	0.4308	0.207	0.2381	-0.0857	-0.1231	-0.1531	-0.2601	0.2946	-0.1163	-0.3856	-0.133
d4	-0.2056	-0.0294	-0.1161	0.2533	0.2738	-0.0934	-0.0294	-0.0422	-0.0525	0.0511	0.0309	-0.0135	0.3018	0.2191

Table 5.7 Correlation table for the remaining traffic, geometric, and environmental factors (continued).

	flowratio	absflowratio	approach_dominance	approach_flow_ph	conflicting_flow_ph	nc2	nc3	nc4	lu2	lu3	lu4	lu5	lu6	d2	d3	d4
app_speed_limit																
b_quad_type																
n_start_lanes																
n_end_lanes																
n_app_lanes																
n_exit_lanes																
n_slip_lane																
b_app_med_type																
b_driveway																
a_quad_size																
r_out_start																
w_appron																
w_lane1																
d_app_inter																
flowratio	1															
absflowratio	0.3878	1														
approach_dominance	1	0.3878	1													
approach_flow_ph	0.5133	0.2043	0.5133	1												
conflicting_flow_ph	-0.4995	-0.3608	-0.4995	0.281	1											
nc2	-0.0171	0.0046	-0.0171	-0.0177	0.0905	1										
nc3	-0.1139	-0.1476	-0.1139	0.3472	0.4946	-0.4714	1									
nc4	0.2045	-0.044	0.2045	0.0415	-0.2969	-0.3939	-0.2476	1								
lu2	-0.108	0.3471	-0.108	-0.1967	-0.2667	-0.0174	-0.2498	-0.3286	1							
lu3	-0.1337	-0.0755	-0.1337	0.0908	0.3672	0.198	-0.0934	-0.078	-0.1239	1						
lu4	-0.2482	-0.2972	-0.2482	0.1625	0.5522	-0.2887	0.7485	-0.2274	-0.3612	-0.0857	1					
lu5	0.036	-0.1026	0.036	-0.0396	-0.0177	0.4714	-0.2222	-0.1857	-0.2949	-0.07	-0.2041	1				
lu6	0.2664	0.1254	0.2664	-0.0078	-0.1402	0.3536	-0.1667	-0.1393	-0.2212	-0.0525	-0.1531	-0.125	1			
d2	-0.1033	0.1686	-0.1033	-0.0533	0.0798	-0.1833	0.3339	-0.5252	0.3823	-0.198	0.2887	0.0236	-0.1473	1		
d3	-0.0144	-0.1235	-0.0144	0.0045	0.076	0.5774	-0.2722	-0.2274	-0.2107	0.343	-0.0714	0.2041	0.3572	-0.5774	1	
d4	-0.1728	-0.0091	-0.1728	-0.0286	0.241	0.198	-0.0934	-0.078	0.2374	-0.0294	-0.0857	-0.07	-0.0525	-0.198	-0.0857	1

Other significant correlations include

- `n_start_lanes` and `n_end_lanes`, which can be explained by the fact that adding or removing lanes is less frequent within the roundabout than at the approach and exit, but still not entirely uncommon.
- `nc3` and `app_speed_limit`, which is unsurprising given that the network class 3 designation is a speed limit of 70 to 90 km/h almost by design.
- `1u5` and `r_out_start`, which is a little surprising, although potentially coincidental and perhaps not explained by a causal relationship, given that the number of roundabouts classified as `1u5` is small (5) and contains a large share of converted traffic circles.
- `1u4` and `nc3`, which can easily be explained by the fact that a large share of the industrial sites are in resource extraction, and the sites are accessed predominantly via regional highways.
- `nc2` and `n_start_lanes`, which can be explained by the fact that a large share of arterial roads host two lanes of traffic in either direction.
- `nc3` and `w_apron`, which is a bit unexpected, but can be explained by the fact that almost all regional highways are managed by the MTQ and this transportation authority has templated implementation of certain design features unrelated to flow directly (apron width is a feature implemented primarily to facilitate the movement of trucks, which are common on regional highways, for the same reason that heavy industry is associated with `nc3`).

The majority of the remaining factors are mostly independent of each other, though a few medium correlations can also be found in Table 5.6 and 5.7. Care will have to be taken during regression when selecting any two of these factors.

5.4 Regression Analysis

A linear regression analysis of speed, yPET, and TTC is performed using geometric and land use factors from Table 5.4. Basic linear regression follows the simple model

$$y_i = \alpha + \sum_k \beta_k X_{ki} + \epsilon_i \quad (5.9)$$

for site $i = 1, \dots, n$ (merging zones), where α is the model intercept, β_k is the coefficient of factor X_{ki} for $k = 1, \dots, m$ factors, and ϵ_i is the regression error. All regressions are subject to normal hypothesis testing, including multicollinearity (detailed in section 5.3.5), homoscedasticity (White’s test), normality of residuals (Shapiro-Wilk W -test), model specification (link test and Ramsey regression specification-error test), and Wald tests of simple and composite linear hypotheses testing.

In the case of TTC_{15} , additional exposure and road user behaviour factors from Tables 4.7 and 5.2 measured at the interaction instant are used. Some of the factors are disaggregated at different reporting levels, and form panel data. Land use factors are typically reported at the roundabout level, while geometric factors are reported at the merging zone level. However, this sampling bias is reduced, since nearly half the roundabouts have only one merging zone under study, . Meanwhile, microscopic exposure, road user behaviour parameters, and TTC_{15} SSMs are reported at the interaction instant level (per road user).

Given that a single merging zone hosts thousands of road users, a special statistical model is required for regressions involving exposure factors: the time-independent, unbalanced, random-effects model, also called a variance components model. This is in contrast to the fixed-effects model, which assumes that every individual component (user pair), contributes to the group equally, and that user pair variance cannot be explained by latent factors. As such, the random-effects model is a special case of the linear fixed-effects model, based on the ordinary least squares estimator. It takes the form of the following equation:

$$y_{ij} = \alpha + \sum_k \beta_k X_{kij} + u_{ij} + \epsilon_{ij} \quad (5.10)$$

for $j = 1, \dots, m$ user pairs and for site $i = 1, \dots, n$ (merging zones), where α is the model intercept, β_k is the coefficient of factor X_{kij} for $k = 1, \dots, m$ factors, u_{ij} is user pair-specific random error (also referred to as the *between* error), and ϵ_{ij} is the “ordinary” regression error (also referred to as the *within* error). Random effects adjust the fixed-effects model with the between-effects model. It models the mean response from means calculated from the user pairs for each group. In this way, the random effects model is a weighted average of the fixed-effects and between-effects models. For this work, the `xtreg` command in Stata 12.0 is used to perform the regression (and more generally the `xt` series of commands and tests, see Appendix C for more information). Random effects modelling using `xtreg` uses a generalized least squares estimator.

In addition to normal hypothesis testing, for each regression, the Breusch and Pagan Lagrangian multiplier test is performed to determine if variance of the unobserved fixed effects

is zero. If this hypothesis is rejected, then ordinary least squares estimator is not appropriate, and the random effects model based on a generalized least squares estimator is warranted.

Additionally, a general Hausman specification test is performed to test the null hypothesis of consistency between estimators of a fixed-effects and random-effects model. Thus, if the null hypothesis is not rejected in a Hausman test, use of random effects is not rejected in favour of fixed effects modelling. Finally, it should be noted that, since random effects regression uses a generalized least squares estimator, R^2 results are similar but not identical to R^2 , calculated from ordinary least squares regression.

Table 5.8 Spearman's Correlation of Merging Zone-Aggregated Surrogate Safety Measure Indicators

	Mean speed	Median Lag yPET	Mean TTC
Mean speed	1.0000		
Median Lag yPET	0.2640	1.0000	
Mean TTC	-0.3242	-0.0769	1.0000

Each of the three SSMs: speed, yPET, and TTC, are only moderately correlated with one another, particularly at an aggregated level, as demonstrated in Table 5.8, suggesting that these indicators are measuring independent behaviours and issues of safety (e.g. speed as a measure of collision severity and TTC as a measure of collision probability). Note that mean statistic is used for speed as it tends to be normally distributed while median is used for yPET as it tends to be heavily-tailed. A mean statistic is used for TTC as it leans towards a normal distribution and tends to have a noisy distribution, especially at sites with limited traffic flow (see Figures 5.17 and 5.18).

At a more disaggregated level, comparison is more tricky, as indicators have differing hierarchies of observation (interaction instant for TTC, user pair for yPET, and interaction instant and road user for speed). Instantaneous TTC is not highly correlated with the instantaneous speed of either road user (-0.1219 and -0.1121, Pearson's linear correlation for the slower and faster road users respectively).

5.4.1 Traffic Parameters

First, a stepwise⁶ linear regression is performed on traffic parameters of hourly traffic volume (`inflow_php1`, measured in veh/h), approach dominance (`approach_dominance`), and absolute

⁶Independent variables are added and removed automatically based on maximum p -values.

flow ratio (`absflowratio`), using all non-collinear geometry and land use parameters from Table 5.4. These models, with coefficients for significant factors, adjusted R^2 , Wald test, and number of observations are provided in Table 5.9. Overall, the hourly traffic volume model has a high goodness-of-fit as measured by R^2 (with an adjusted $R^2 = 0.707$), while the two flow ratio models score only moderate prediction power (adjusted $R^2 = 0.512$ and $R^2 = 0.431$).

Hourly traffic volume can be explained positively by many factors, including commercial and industrial land use, and increasing urban density; it can be explained negatively by residential land use (as opposed to no land use). Note that hourly traffic volume (`inflow_phpl`) already normalizes for number of lanes. Thus, its positive association with number of start lanes suggests that roundabout merging zone capacity is more fully realized at multi-lane facilities than at single-lane facilities, but that multi-lane approaches and exit capacity are underutilized at multi-lane facilities. Painted medians and increasing upstream distances are also positively associated with hourly traffic volumes, however these effects might be better explained by planned implementation as a *result* of higher *planned* traffic demand rather than direct inducement of demand.

Since demand is correlated with number of lanes, and demand dictates resulting flow ratios, approach dominance can, unsurprisingly, be heavily explained by the number and configuration of approach and exit lanes. The opposite is true of number and configuration of start lanes, for similar reasons. The slip lane is positively associated with approach dominance; its implementation is almost surely the result of excessive demand on the approach. Note that slip lanes are relatively uncommon. Of the 35 merging zones studied, only one had a slip lane. Consequently, although statistical models show statistical significance, slip lane results should be interpreted cautiously, or rejected entirely. Roundabout merging zones serving as access ramps to and from limited-access-highways are positively associated with approach dominance, possibly because these highway ramp flows govern overall demand at the roundabout (undevelopped land). It should be noted that, in this study, roundabouts predominantly serve as highway exits, and therefore also as roundabout approaches). Similar trends can be seen for commercial and industrial roundabouts serving points of interest.

The absolute flow ratio, a measure of traffic stream polarization, has slightly fewer significant factors. Factors predicting polarization positively include number of exit lanes and low density. Meanwhile, industrial land use and number of slip lanes tend to predict a decrease in polarization.

In general, parameters of traffic and exposure tend to be highly correlated with geometric factors—not so surprising given that engineers make design choices based on expected

Table 5.9 Linear Regression Models for Exposure

Model	inflow_phpl		approach_dominance		absflowratio	
	Coefficient	P> t	Coefficient	P> t	Coefficient	P> t
_cons	447.37	0.000	-1.093	0.008	0.624	0.021
app_speed_limit	-3.025	0.014	-	-	-	-
a_quad_size	-	-	-	-	-0.00374	0.056
b_app_med_type	175.81	0.000	-	-	-	-
b_quad_type	-195.16	0.010	0.590	0.006	0.384	0.083
d_app_inter	0.104	0.005	0.000530	0.001	-	-
d2	76.71	0.010	0.814	0.001	0.150	0.058
d3	-	-	0.936	0.000	-	-
d4	389.83	0.000	-	-	-	-
lu2	-147.05	0.000	-	-	-	-
lu3	316.52	0.000	-0.426	0.030	-	-
lu4	139.86	0.001	-0.167	0.049	-0.225	0.018
n_app_lanes	-159.39	0.001	0.293	0.003	-	-
n_exit_lanes	-97.17	0.004	0.295	0.006	0.256	0.006
n_slip_lane	-	-	0.432	0.039	-0.568	0.014
n_start_lanes	87.23	0.018	-0.245	0.059	-	-
nc4	-	-	0.740	0.000	-	-
r_out_start	-	-	-	-	-0.0106	0.022
w_lane1	-	-	0.046	0.061	-	-
Adjusted R^2	0.7070		0.5123		0.4307	
Wald prob. > F	0.0000		0.0025		0.0015	
Observations	35		35		35	

demand. As such, the remaining regression models generally can not include both traffic and geometric factors. This concept—that exposure, collision rates, and geometry are all interlinked with one another—has been explored before (e.g. in Ewing and Dumbaugh, 2009; Miranda-Moreno et al., 2011b; Strauss et al., 2013). It follows that geometry explains traffic (including exposure) and traffic explains road safety to some degree, and this association is the hypothesis that is made in the following analysis.

5.4.2 Speed

A stepwise linear regression is performed on mean road user merging zone speed (measured in km/h) for all road users (`mean_speed`), using first geometry and land use parameters from Table 5.4, then traffic parameters plus unused geometric factors from the traffic parameter regression in section 5.4.1. These models, with coefficients for significant factors, adjusted R^2 , Wald test, and number of observations are provided in Table 5.10.

Table 5.10 Linear Regression Model for Mean Road User Speed

Model	Geometry-only model		Traffic model	
	Coefficient	P> t	Coefficient	P> t
_cons	27.69	0.000	27.815	0.000
app_speed_limit	.1222	0.022	-	-
b_quad_type	20.02	0.000	-	-
d2	3.017	0.030	-	-
n_slip_lane	-9.135	0.024	-	-
w_apron	-1.196	0.005	-	-
approach_dominance	-	-	7.355	0.079
Adjusted R^2	0.5396		0.0629	
Wald prob. > F	0.0000		0.0000	
Observations	35		35	

Regressing only for geometric and land use factors leads to a moderately predictive model, with an adjusted $R^2 = 0.540$. Significant factors associated with increases in mean merging zone speed include

- an increase in speed limit, though by only a tenth of the rate, corroborating the school of thought that holds that posted speed limits have only a marginal effect on modifying road user speeds;
- irregular merging zone shape, or more specifically, one approach serving two exits;
- medium urban density;
- lack of a slip lane, though the speed on the slip lanes is not captured; it is possible that movement type (iv) road users travel faster and are simply not captured in this sample, and furthermore, as stated earlier, caution is warranted when interpreting this factor, as the number of samples is very low; and
- shorter apron widths.

The traffic-parameters-only model provides poor predictive power, but it does suggest that approach dominance is positively correlated with mean speed road user merging zone, affecting up to just over 7 km/h.

Shorter apron widths are known, through experimental observations and inspection of trajectory maps, for causing issues with the reduced deflection of road users, and thus straighter and faster through-movements by road users.

5.4.3 Yielding Post-Encroachment Time

Next, a stepwise linear regression is performed on median lead and lag yPET below 5 s at each site. Recall that yPET is not very normally distributed (it is at the very least a mixture model with a normal and some negative exponential components), so median aggregation is preferred.

Median lead yPET below 5 s for approach road users could not be explained by any geometry, land use, or traffic parameter, except for hourly traffic volume (for a model using only `inflow_ph_p1` as an independent variable), and only barely, with an adjusted $R^2 = 0.0780$, and a coefficient of 0.0027 with marginal statistical significance. It seems that the tendency for approaching road users to follow conflicting flow road users with a median headway of 2 s remains consistent regardless of any external factors, except when traffic is low enough that following road user behaviour decreases (arrivals are correspondingly less uniform, given fewer queues) .

Median lag yPETs below 5 s are a different story however. No discernible pattern exists between sites in lag yPET measures, as seen in section 5.3.3. However, at individual sites, it seems that lag yPETs can be explained by merging zone characteristics. A stepwise linear regression is performed on median lag yPETs below 5 s across all merging zones. These results are shown in Table 5.11. Recall that the smaller the yPET measure is, the closer road users make their trajectory negotiations: more specifically, lag yPET is the time an approach road user has before the next conflicting road user, which always has priority, enters the same space, such as a merging zone. Thus yPET serves as a proxy to yielding, and collision probability.

Table 5.11 Linear Regression Model for Median Lag yPET below 5 s

	Coefficient	P> t
<code>_cons</code>	1.657	0.000
<code>app_speed_limit</code>	-0.0137	0.037
<code>b_driveway</code>	-1.022	0.003
<code>d4</code>	-1.276	0.037
<code>nc4</code>	-0.579	0.034
<code>w_lane1</code>	0.203	0.002
Adjusted R^2	0.4085	
Wald prob. > F	0.0011	
Observations	35	

From these results, it can be deduced that significant reductions in yPET (to below 5 s), and hence hypothetical increases in collision probability according to surrogate safety method

theory, is associated with

- an increase in the speed limit (e.g. a 10 km/h increase yields a 0.1 s median yPET decrease), though it is interesting to note that the measured road user mean speed is *not* significantly associated with lag yPET below 5 s (possibly explained as the speed limit being a proxy for other factors);
- presence of a driveway, within or in the immediate vicinity of the merging zone, a practice that should be avoided, even if only to reduce road user trajectory negotiations;
- high urban density;
- limited-access-highway ramps, possibly explained by road users exiting the highway and who have not quite transitioned into non-highway behaviour;
- a decrease in lane width. This is somewhat contradictory to what is expected: as lane width decreases, lane sharing, i.e. being within the same designated lane, is expected to decrease (see section 5.3.1), resulting in more uniform arrivals. More investigation is needed with regard to this.

5.4.4 Time-to-Collision

Recalling the framework laid out in Chapter 4, two types of analyses are available: SCC, which attempts to measure a uniform increase in proportion of large TTC_{15} safety indicators; and SEC, which minimizes the expected number of serious events, defined by a threshold $\zeta = 1.5$ s on individual TTC_{15} values, retaining the values below the threshold. Since the TTC_{15} safety indicators used for this analysis are generated from collision courses predicted using discretized motion pattern motion prediction, collision course probabilities are available, and the serious events can be weighed, as an expected number of serious events wSE , by the collision course probability $P(CP)$.

Safety Continuum Comparison

A SCC random effects regression of motion pattern-based 15th centile TTC (measured in seconds) is performed using two models: approach dominance with non-collinear geometric factors, and a traffic parameters-only model (since most geometric parameters are captured by at least one of these traffic parameters and the remaining factors are not found to be significant).

To reflect the earlier probability distribution curve fitting, the log of TTC_{15} is used for this regression instead of TTC values directly. Thus, the random effects regression model takes the shape of

$$\ln(TTC_{15ij}) = \alpha + \sum_k \beta_k X_{kij} + u_{ij} + \epsilon_{ij} \quad (5.11)$$

using similar parameters as in equation 5.10.

Modelling results are shown in Table 5.12. Both models are reasonably good predictors of TTC_{15} , especially between panels (merging zone factors). The within (individual road user) effects offer mediocre predictive power, but suggest that TTC_{15} increases—and therefore collision probability hypothetically decreases—with increased 15 s exposure (i.e. the “safety in numbers” effect). Recall that the instant when speed, exposure, angle, etc. are measured corresponds to the instant when the 15th centile collision course occurs.

The model also shows an increase in TTC as the angle of incidence increases, i.e. higher for side-swipe conflict-type interactions than for rear-end conflict-type interactions. This may be explained by the fact that *either* road user has the opportunity to see the other in such a situation, while rear-end conflicts may be more unexpected for the leading road user. Roundabouts work because of the interactive behaviour of road users as they approach the merging zone, generating many more side-swipe conflict than road users may be used to.

Increased speed of the *slower* road user (`lower_inst_speed`) at the interaction instant of the collision course is found to be associated with lower TTC. The same cannot be said for the speed of the *faster* road user; however, this does not suggest that the differential velocity does not play a factor. Recall equations 4.9 and 4.11 in the calculation of TTC. It should be noted that the speed of either road user is found to be moderately correlated (0.695). This is not surprising, given that some degree of homogeneity should be expected of similar road users in the same environment, e.g. motorists in a roundabout with a single posted speed limit, but that some variation should exist too, e.g. variation in individual road user characteristics and different yielding rules between lanes.

The between effects are as follows:

- An increase in merging zone size (`a_quad_size`) and roundabout radius (`r_out_start`) is associated with increases in TTC.
- Presence of a driveway (`b_driveway`), an irregular merging zone design (`b_quad_type`), and medium urban density (`d2`) are associated with a decrease of TTC. The first two parameters are non-standard design features and should be avoided. Urban density

may reflect generally increased road user activity and thus interaction complexity.

- In the traffic model, it can be seen that when traffic flow favours the approach, or, alternatively, when traffic flow is more balanced between the approach and merging zone start (`absflowratio`), TTC increases. It is possible that when traffic flow is very unbalanced towards the conflicting flow, approaching road users must wait longer (yield) for a gap and may take more risks. It is worth reminding that `approach_dominance` and `absflowratio` are modestly correlated (0.38), meaning that there is a small amount of overlap in the results between these two factors.

Note that mixing the factors in the two models in a third, yields a between R^2 in excess of 0.799. It indicates potential for good predictive power, and suggests that neither model overlaps the other fully.

Table 5.12 Random-Effect Regression Models for Motion Pattern-Based 15th Centile Time-to-Collision (SCC)

Model	Global model		Traffic model	
	Coefficient	P> t	Coefficient	P> t
<code>_cons</code>	-	-	0.8326	0.000
<code>a_quad_size</code>	0.0063	0.012	-	-
<code>b_driveway</code>	-0.6915	0.000	-	-
<code>b_quad_type</code>	-0.6868	0.033	-	-
<code>d2</code>	-0.1864	0.069	-	-
<code>r_out_start</code>	0.0132	0.022	-	-
<code>approach_dominance</code>	0.5006	0.028	0.8154	0.002
<code>absflowratio</code>	-	-	-0.5930	0.007
<code>fifteen_second_exposure</code>	0.0151	0.000	0.0151	0.000
<code>interinst_angle</code>	0.0029	0.000	0.0029	0.000
<code>lower_inst_speed</code>	-0.0179	0.000	-0.0179	0.000
Within R^2	0.1245		0.1245	
Between R^2	0.6709		0.4220	
Overall R^2	0.2515		0.1338	
Wald prob. > F	0.0000		0.0000	
Observations	32111		32111	
Groups	35		35	

Serious Event Comparison

A stepwise linear regression of discretized-motion-pattern-based serious TTC events (SE_{TTC} measured in events per hour) is also performed, using weighted SEC methodology with a

threshold of $\zeta = 1.5$ seconds. Among all the geometry, land use, and traffic parameters available and studied up to this point, only a single factor is found to have any statistical significance in modelling serious TTC events: traffic volume. Given that the SEC methodology is essentially a count of number of events, the correlation between it and traffic volume is not entirely surprising, given that traffic volume produces traffic events, or user pairs. However, some degree of variation between sites, after controlling for exposure, would ordinarily be expected. Given the lack of any other factors yielding significant SEC prediction, the usefulness of SEC questionable.

Table 5.13 Linear Regression Models for Motion Pattern-Based Time-to-Collision Events Below 1.5 s Per Hour (Weighted SEC)

	Coefficient	P> t
<code>_cons</code>	-2.144	0.322
<code>inflow_phpl</code>	0.02839	0.022
Adjusted R^2	0.1236	
Wald prob. > F	0.0218	
Observations	35	

5.5 Summary

In this chapter, an application of the automated video-based traffic data collection system and expanded surrogate safety method framework is presented, using roundabouts in the province of Québec as a case study. It was demonstrated that roundabout geometry and land use are reflected in traffic parameters, and that these traffic parameters are in turn reflected in the various aspects of road safety (summarized in Figure 5.21) via several SSMS, and that these SSMS are mostly independent of one another. The preceding observation is consistent with observations regarding the relationship between built environment, traffic parameters, and road safety made in the literature (Ewing and Dumbaugh, 2009; Miranda-Moreno et al., 2011b; Strauss et al., 2013).

Regression models were prepared in order to explain traffic parameters (inflows, flow ratios, etc.) and SSMS from geometry and land use adequately. A number of factors were found to be associated with one or more elements of road safety. Generally speaking, it is found that small aprons (consequently generating poor road user deflection at the approach) result in higher observed road user speeds. Increasing posted speed limits is found to be associated with proportional increases in speed, though only at a rate of one tenth the posted speed

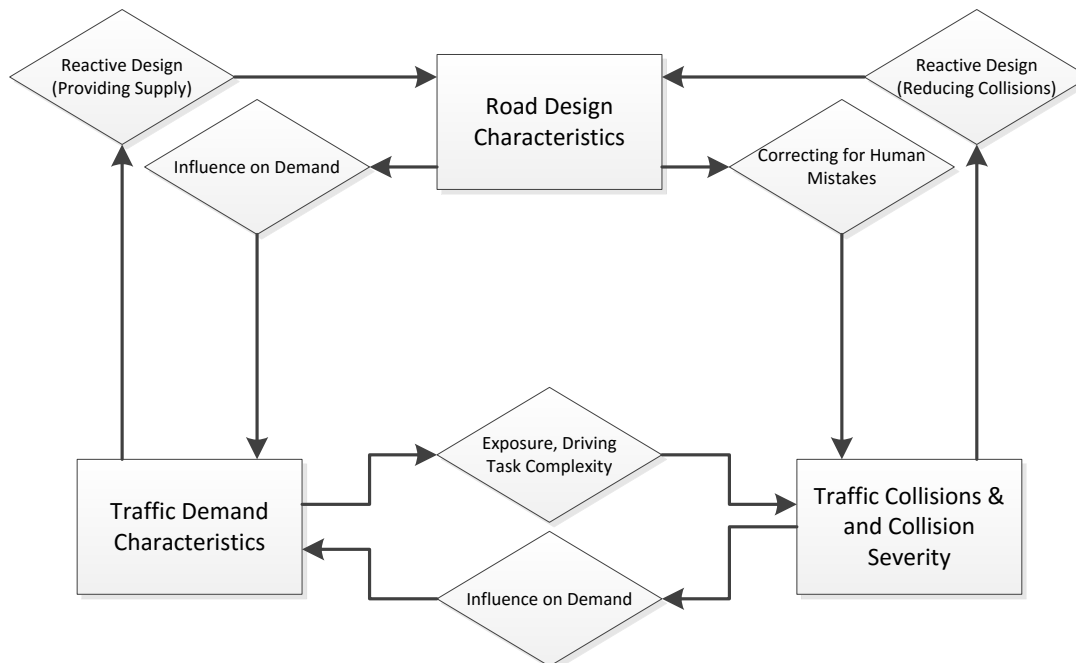


Figure 5.21 Relationships between traffic characteristics, road design, and traffic safety.

limit. Designs which deviate from typical roundabout design, such as irregular merging zone configuration, are found to be associated with higher speeds and lower TTC measures, while the presence of driveways on or immediately near merging zones is associated with reduced yPET and TTC measures. Consequently, these designs should be avoided. In the case of one-way roads attached to a roundabout causing irregular merging zones, merging issues might be mitigated with larger merging zones and a more homogenized flow. The regression models are summarized in Table 5.14. The effect of lane width on yPET and presence of slip lanes is somewhat uncertain and will require further investigation.

Furthermore, while simple conversion of traffic circles into roundabouts (without effecting comprehensive geometry changes) was not parametrized in these models, this factor was found, through cluster analysis, to be associated with poor TTC safety indicators. As a result, the recommendation of avoiding this practice in general is made, as the original physical dimensions of the traffic circle rarely meet the design requirements of the basic roundabout design.

Finally, following the results presented in Chapter 4, it is shown experimentally that SEC is of limited use for investigating potential contributing factors of road design to aspects of road safety in this study. Given the general issues outlined with SEC, its use is discouraged. Many more factors are correlated with and can explain speed, lag yPET, and TTC, demonstrating

Table 5.14 Regression Model Summary

Model	Equation	R^2	Units
Speed	$= 27.7 + 0.122V_{limit} + 10.0b_{i_{merge}} + 3.0d_2 - 9.1n_{slip} - 1.2w_{apron}$	0.5396	km/h
Lag yPET	$= 1.65 - 0.0137V_{limit} - 1.02b_{driveway} - 1.28d_4 - 0.58nc_4 - 0.20w_{lane}$	0.4085	s
Time-to-Collision	$= \overline{TTC} + e^{c_{TTC}}$	0.6709	s
Expected Number of Serious Events	$c_{TTC} = 0.0063a_{merge} - 0.69b_{driveway} - 0.34b_{i_{merge}} - 0.19d_2 + 0.013r_{outside} + 0.5app_{dom} + 0.0151exp_{15s} + 0.0029\overline{a_{int}} - 0.018\overline{V_{lower}}$ $= -2.144 + 0.02839Q_{lane}$	0.1236	events/h

the usefulness of SCC-based surrogate safety methods.

Use of the merging zone as a unit of analysis for roundabouts is novel, given the microscopic nature of the surrogate safety methods used. However, its use does increase complexity for comparison with non-roundabout intersections.

In the next chapter, a similar international analysis of roundabout merging zones is performed to determine if aspects of road safety can also be explained from elements of road user behaviour and local cultural factors regarding road use.

CHAPTER 6 INTERNATIONAL COMPARISON BETWEEN QUÉBEC AND SWEDISH ROAD USERS AT ROUNDABOUTS

While the broad concepts behind road design and signalization are universally recognized for the sake of road user mobility between regions of the world—e.g. to accommodate visitors—specifics of intersection design philosophy and signalization differ significantly between North America and Europe. This is not surprising, given that the United States and Canada are not signatories of either the *1949 Geneva Protocol on Road Signs and Signals* or the *1968 Vienna Convention on Road Signs and Signals* which codify road signalization throughout nearly all of Europe and much of Asia. Instead, intersection design in the United States and Canada, and much of the Pacific, is on the *1935 Manual on Uniform Traffic Control Devices*.¹

Design differences are especially notable with level I and level II signal control at intersections, i.e. intersections without traffic lights. While European design tends to favour a limited use of stop signs in favour of yield signs or implicit priority for right-approaching road users, i.e. no signs at all, North American design favours stop signs almost exclusively. In fact, yield signs in North America are generally used for slip lanes or merging zones only, and never to control intersections directly. Given the yielding nature of the roundabout design, it is not surprising that roundabout adoption has been very late in North America. While roundabouts are a relatively new phenomenon in North America, they have existed since 1966 in the United Kingdom where the modern design of the roundabout was first conceived.

However, roundabouts are beginning to flourish in North America, against the prevailing stop-sign-predominant intersection design philosophy. Thus, studying this discrepancy in road design philosophy is especially relevant today, as many North American road users may not be familiar with the non-stop-controlled intersection design that roundabouts introduce; this has been cited in the literature as a problem to short-term and medium-term adoption of the roundabout (e.g. Retting et al., 2007). To this end, there exists a need to study differences in driving culture between the two continents, whether induced or latent. In this chapter, an international, microscopic comparison of road user behaviour between road users in Québec and Swedish roundabouts is made. While many international studies of road safety and design have been conducted, to date, few behavioural comparison studies have been attempted, and none have been conducted at the level of detail, and on the scale, of this study.

For this study, a video data set complementary to the data collected in Québec for Chap-

¹The latest edition dates from 2009 with updates in 2012.

ter 5 is collected at a number of Swedish roundabouts comparable to those being studied in Québec. The analysis procedure for roundabouts is largely the same as for that of roundabouts conducted in Chapter 5, although particular attention is paid to comparative analysis between the two regions, and particularly to measures of road user behaviour. A discussion and comparison of historical accident data, available at a more macroscopic level, is also included to provide context for the surrogate safety methods.

6.1 Rationale

Sweden and Canada are among some of the safest countries in the world for motorists, cyclists, and pedestrians, by a wide margin. Despite this, annual traffic fatality rates in Canada are nearly twice as high as in Sweden, as measured per 100,000 inhabitants (World Health Organization, 2013; OECD and ITF, 2015)), per 10,000 registered motor vehicles (OECD and ITF, 2015), and per billion veh-km travelled (OECD and ITF, 2015), and this despite a relatively comparable car occupancy rate and mode share (OECD and ITF, 2015). Figures of reported accidents per 100,000 inhabitants share a similar trend (OECD, 2015). These numbers are summarized in Table 6.1. While all these rates have been observed to be decreasing consistently over the last 40 years, the difference between the two countries has been fairly constant (OECD, 2015). Furthermore, fatality and accident rates in Québec are consistent with, and thus representative of, the Canadian national average (Transport Canada, 2015).

Table 6.1 Comparison of Macroscopic Historical Accident Statistics

Statistic	Canada	Sweden	Ratio	Year	Source
Fatalities per 100,000 inhabitants	6.8	3.0	2.26	2010	(World Health Organization, 2013)
Fatalities per 100,000 inhabitants	5.5	2.7	2.03	2013	(OECD and ITF, 2015)
Fatalities per 10,000 registered motor vehicles	0.85	0.45	1.88	2013	(OECD and ITF, 2015)
Fatalities per billion veh-km travelled	5.6	3.4	1.65	2013	(OECD and ITF, 2015)
Accidents per 100,000 inhabitants	501.0	209.8	2.39	2013	(OECD, 2015)

This disparity in road safety between Sweden and Québec is unexplained, given that both

share similar population sizes, levels of urbanization² (81 % in Québec versus 86 % in Sweden), climate, and economic development factors. Instead, the disparity might be explained by one of two sets of microscopic factors: road design, or road user behaviour and road culture (Moeckli and Lee, 2007; Ward et al., 2014). To this end, a comparative study of road user behaviour between the two regions is prepared by first isolating geometric factors, as these are the most straightforward to control. Roundabouts are arguably among the best types of road infrastructure for direct comparison between Canada and Sweden. Although a relatively new phenomenon in North America, roundabouts are among the few types of road design that can be found with near identical geometric and aesthetic features in both regions of the world. This should come as no surprise, as the North American roundabout design guides (Rodegerdts et al., 2010) are heavily influenced by European roundabout design. One important difference is that cycling is significantly more prevalent in Sweden than in Canada or Québec. Fortunately, roundabout design in Sweden tends to favour grade separated cycle crossings (Sakshaug et al., 2010). All the roundabouts sampled for study in this chapter have either a negligible cycling demand, or reroute the traffic elsewhere.

Sweden is chosen as a candidate European region to compare with Québec for roundabouts due to the abundance of roundabouts in the region and because, as stated earlier, these regions share similar demographics, levels of economic development, and climate characteristics. Researchers from Lund, Sweden, collaborated on this project, as they had aligned interests in advancing and applying surrogate safety methods, particularly TCTs (Hydén, 1987), but also more modern approaches with the use of automated video-based data collection systems (Laureshyn et al., 2009, 2010; Laureshyn, 2010; Svensson et al., 2011). This same research group at Lund University had already produced early road user behaviour and road safety analyses using a variety of surrogate safety methods to investigate roundabouts in Sweden, in particular cyclist crossings (Sakshaug et al., 2010) and roundabout conversions in the Swedish city of Växjö (Hydén and Várhelyi, 2000). Other relevant roundabout work in Lund includes a study of multi-lane roundabout gap measures using some manually-annotated video data to calibrate microsimulation software (Irvenå and Randahl, 2010).

Despite the abundance of roundabouts in Sweden, the roundabouts sampled in Lund have been implemented much more recently, many in the last two decades, than in other areas of Europe, such as the United Kingdom, as early as 1966, and France. The median date of construction of roundabouts in the province of Québec is 2008 as of writing. The median age of the Swedish roundabouts sampled for this study is 12 years, compared to 9 years for the Québec sites. Construction years are available in Tables 6.3 and 6.4, while data collection

²The populations of Canada and Sweden tend to be concentrated at the southern edges and along the coasts of each country.

took place in 2016 and 2013 in Sweden and Québec respectively.

6.2 Swedish Roundabout Data Inventory

An inventory of 136 roundabouts in and near the cities of Malmö and Lund in the Skåne county of Sweden is created to mirror the inventory created for the Québec safety project as discussed in section 5.2. The database is created partially using the GoogleMaps API to pull GIS data automatically from identified roundabout coordinates. Additional geometric factors are added as needed to reflect most of the relevant geometric factors examined in Chapter 5.

6.2.1 Site Selection

Figure 6.1 maps the location of the 136 roundabouts in the roundabout inventory prepared for the Malmö and Lund regions (as of 2015) of Skåne County, the southernmost region of Sweden. The roundabouts targeted for analysis with video data collection are highlighted. They are selected using a ten-point scale similar to that used for Québec roundabouts, with weighted averaged scores for multiple factors including feasibility of performing video data collection, and geometric and built-environment similarity with the roundabouts studied in Québec (see Appendix D for details). Ultimately, time and budget were allocated for data collection at four roundabouts, all in the periphery of the city of Lund. The VCU deployed for this data collection differs significantly from the one used in Québec: instead of attaching to existing street poles, it uses a trailer with a hoist capable of lifting two cameras to a height greater than 15 meters. Typical installation is illustrated in Figure 6.2. The VCU is deployed at four roundabouts, generating a total of eight camera views, each covering one or more roundabout merging zones. Each camera view collects four to five days of traffic data, from 6:00 a.m. to 10:00 p.m.




Figure 6.1 Roundabout locations across Skåne county, Sweden.

6.2.2 Data Inventory

Table 6.2 summarizes the size of the data collected for the Swedish sites. Fewer sites are visited in total than in Québec, but each site has data collected for more merging zones, and for a longer period of time. This reflects the more permanent nature of the VCU used in Sweden for data collection. Overall, the size of the Swedish data set is about 20 % the size of the Québec data set collected as part of the methodology discussed in Chapter 5, however, only a sub-sample of comparable size is used from Chapter 5 to represent Québec roundabouts in the study in this Chapter. These Québec sites are selected to match the sites in Sweden and *vice versa*.

Note that the figure in Table 6.2 representing the number of hours analysed is greater than the total number of hours of video data captured. This is because some of the camera views covered more than one merging zone, creating multiple analysis zones at each. Furthermore, because of this and because all the sites are clustered in Lund, it is not reasonable to assume that the road users captured in the video data are all unique, unlike in the Québec data. Given the population and size of Lund and the distribution of the four roundabouts throughout the city, it is estimated that between 5,000 to 15,000 unique Lund road users are captured as part of the study. Given that all observations are performed for road users in Lund only, use of this sample as a representative sample of Sweden requires the assumption that road user behaviour is uniform in the region. This limits the conclusions somewhat, though it is reasonable to assume that cultural factors are much more homogeneous in subcontinental regions as small as Sweden than between continents, especially considering that road users are not stationary, neither are they bound to individual cities; instead, road users typically operate within a territory that includes inter-city movement.

The Québec and Swedish roundabouts selected for this study are selected on the basis of similarity of geometric design and land use, and close resemblance to the most typical implementation in either region. This configuration is demonstrated in Figure 6.3. It is characterized by

- single lane configurations (on each of the approach, exit, and roundabout ring);
- an approach speed limit of 50 km/h;
- an outside radius of 15-25 metres;
- very-low to medium urban density;

Table 6.2 Sweden Roundabout Data Inventory

Data collection		
roundabouts visited	4	
Camera views	8	
Merging zones/analysis zones	10	
Total hours of video (pre-analysis)	275.3 h	
Disk space usage		
Video data	360,626 MB	(570 files)
Trajectory data	68,771 MB	(570 files)
Cached data	27,211 MB	(1,209 files)
Misc. overhead (e.g. system)	5,193 MB	(2,993 files)
Trajectory Data		
Unique road users	66,513	
Vehicle-kilometres travelled	8,107.7 veh-km	
Duration (analysis)	332.6 h	
Interaction Data		
User pairs	137,111	
User pairs with TTC (cmp , 15t)	18,739	

- suburban, commercial, or mixed land use;
- moderate visibility over the centre island (road users on opposing merging zones are only partially obstructed from view);
- and presence of pedestrian crosswalks on all or nearly all approaches and exits.

Some of the Québec roundabouts had a 35 km/h speed advisory posted at the roundabout approach as well. Swedish roundabouts do not have posted speed advisories (Isebrands, 2011).

As with the previous study of roundabouts in Chapter 5, roundabout merging zones are used as the unit of analysis. While some roundabout observations are repeated two to four times amongst the merging zones, it is very important to note that the road users and their interactions are independent from one merging zone to the next, even within the same roundabout, given the slight differences in geometry, significant differences in traffic conditions, and, more importantly, given that different merging zones within even the same roundabout serve mostly different movements of traffic (recall the merging zone movements defined in section 5.1.2). The exception to this rule is for adjacent merging zones serving predominantly right-exiting traffic, i.e. movements (iii), (i), and (ii), in that order.

Table 6.3 lists each of the Swedish sites and merging zones examined, and Table 6.4 lists

each of the Québec sites and merging zones examined. Some of the more distinguishing features (even after control of geometry and land use, no two sites are perfectly identical) are summarized and averaged in each table. The average age of the Swedish roundabouts, and thus merging zones, sampled in this study is 20 years, though there is significant variance, ranging from 5 years to 50 years. Historical accident data, pulled from national records that sample a 2 to 15 year period at each roundabout, average 1.8 accidents *p.a.*. Historical accident data, pulled from the sites sampled in Québec as part of the Québec Roundabout Safety research project (Saunier et al., 2015), are more recent on average, at 10 years, however, none as recent as the most recent Swedish site sampled. These sample a period covering 4 to 7 years at each roundabout and average 4.19 accidents *p.a.*. Admittedly, these historical accident data suffer the very same issues that hamper historical accident analysis, justifying the development of surrogate safety methods. However, these historical accident data are consistent mostly with national averages of either region and between regions, suggesting that the sites sampled are representative of differences in the macroscopic (national-level) road safety observations noted in the discussion earlier.



Figure 6.2 Example installation of a trailer-type video data collection unit.



Figure 6.3 Sample camera view of a roundabout merging zone, illustrating typical roundabout design features. The site depicted is SE-Svenshogs-NorraGrans-2.

Table 6.3 Site Characteristic Inventory: Sweden

Site	Land Use	Urban Density	Outside Radius (m)	Hourly Flow (veh/h/ln)	Flow Ratio	Construction Year	Accidents <i>p.a.</i>
SE-Fasanvagen-Trollebergsvagen-1	Mixed	Medium	25.0	408.4	-0.432	1965	4.1
SE-Fasanvagen-Trollebergsvagen-2	Mixed	Medium	25.0	394.7	0.283	1965	4.1
SE-R103-Foretagsvagen-1	Mixed	Very low	22.0	281.8	0.293	2003	1.0
SE-R103-Foretagsvagen-2	Mixed	Very low	22.0	289.0	0.517	2003	1.0
SE-R103-Foretagsvagen-3	Mixed	Very low	22.0	226.8	0.252	2003	1.0
SE-R103-Foretagsvagen-4	Mixed	Very low	22.0	218.4	0.934	2003	1.0
SE-RubenRausings-Borgs-1	Mixed	Low	22.0	123.9	0.646	2010	1.5
SE-RubenRausings-Borgs-2	Mixed	Low	22.0	121.4	0.568	2010	1.5
SE-Svenshogs-NorraGrans-1	Residential	Low	16.5	191.0	-0.417	1995	1.4
SE-Svenshogs-NorraGrans-2	Residential	Low	16.5	142.9	0.054	1995	1.4
MEAN			21.5	239.9	0.270	1995	1.8

Table 6.4 Site Characteristic Inventory: Québec

Site	Land Use	Urban Density	Outside Radius (m)	Hourly Flow (veh/h/ln)	Flow Ratio	Construction Year	Accidents <i>p.a.</i>
QC-desSoeurs-duGolf	Residential	Medium	25.0	315.1	-0.327	2004	7.0
QC-desSoeurs-ReneLevesque	Residential	Low	22.5	178.8	0.421	2003	1.4
QC-Frechette-AnneLeSeigneur	Mixed	Low	24.5	51.5	0.600	2003	7.0
QC-desSources-Riverdale	Residential	Low	18.5	236.9	-0.361	2003	0.7
QC-Mouettes-Alouettes-1	Residential	Low	15.5	64.6	-0.518	2004*	
QC-Mouettes-Alouettes-2	Residential	Low	15.5	93.4	0.607	2004*	
QC-StEmilie-StDenis	Residential	Low	18.5	46.6	0.112	2005	1.0
QC-Talbot-JacquesCartier-1	Mixed	Medium	18.0	150.8	0.608	2004	7.7
QC-Talbot-JacquesCartier-2	Mixed	Medium	18.0	238.6	0.534	2004	7.7
MEAN			19.5	152.3	0.186	2004	4.19

* No construction date available. Date is estimated from historical aerial footage and carries an uncertainty of ± 2 years.

6.2.3 Potential Contributing Factors

As mentioned earlier, most potentially contributing factors outlined in section 5.2.3 have already been controlled for, except for some of the minor discrepancies in radius and land use highlighted in Tables 6.3 and 6.4. In addition to these and the traffic-flow parameters outlined in Table 5.2 (absolute flow ratio, approach dominance, traffic volumes, microscopic exposure, etc.), a small number of potentially contributing factors are added to the study exclusively for international comparison, including years since built, and presence of a crosswalk at the approach and/or exit. These are summarised in Table 6.5.

Table 6.5 Additional Factors

Variable (merging zone)	Factor description	Type (Units)
approachcrosswalk	Presence of pedestrian crosswalk on roundabout approach	Categorical
exitcrosswalk	Presence of pedestrian crosswalk on roundabout exit	Categorical
Variable (roundabout)	Factor description	Type (Units)
yearsSinceBuilt	Years between construction of roundabout and collection of behavioural data	Numerical (Years)

While not a contributing factor, historical accident data is also available for each roundabout and could potentially serve as a dependent variable for regression with SSMs for the sake of calibrating surrogate safety method, if the sample size was more significant and if these historical accident data could be sampled more accurately at the merging zone level. This variable is summarized in Table 6.6.

Table 6.6 Accident Data

Variable (roundabout)	Factor description	Type (Units)
accidents	Yearly accidents as reported by local ambulance services (Sweden) or police services (Québec)	Numerical (accidents/year)

The distribution of median angle of incidence per user pair, as calculated following section 5.3.4, is presented in Figure 6.4. The general shape is roughly comparable to that of the




Figure 6.4 Distribution of median angle of incidence per user pair.

distribution in Figure 5.19, though with an important difference: in increased proportion of road-user-following. Figure 5.19 also shows very little variation between Québec and Swedish merging zones.

6.2.4 Analysis of Correlation of Factors

As was the case with the study of Québec roundabouts in the previous chapter, an analysis of Correlation is performed on the potentially contributing factors before regression. However, since the purpose of this study is to predict aspects of road safety between Swedish sites (`sweden= 1`) and Québec sites (`sweden= 0`), collinearity between factors other than the `sweden` variable are less important. The results of Pearson's linear correlation are presented in Table 6.7. Spearman's correlation is roughly identical, with only the relationships between `yearssincebuilt` and `d3`, and `yearssincebuilt` and `approachcrosswalk` differing by more than 0.1, and neither have a correlation greater than 0.5 with either method.

Table 6.7 Pearson's linear correlation table for the traffic, geometric, and environmental factors

	sweden	lu2	d2	d3	insideRadius_m	outsideRadius_m	approachcrosswalk	exitcrosswalk	yearsSinceBuilt	accidents	hourlyflowvehhln	inflowphpl	absflowratio
sweden	1												
lu2	0.3825	1											
d2	-0.169	-0.5367	1										
d3	-0.2469	0.2066	-0.6086	1									
insideRadius_m	0.1804	0.2106	-0.3037	0.0974	1								
outsideRadius_m	0.1352	0.4115	-0.3467	0.2311	0.7897	1							
approachcrosswalk	-0.0738	-0.019	0.1273	0.2988	0.0438	0.0556	1						
exitcrosswalk	-0.3055	-0.2697	0.3443	0.2357	-0.1664	-0.1052	0.3099	1					
yearsSinceBuilt	0.3919	0.1475	-0.3192	0.4854	0.3132	0.3406	0.145	0.096	1				
accidents	-0.5355	0.232	-0.3547	0.7781	-0.0935	0.1767	0.3668	0.2754	0.0824	1			
inflowphpl	0.32	0.201	-0.7091	0.5663	0.3212	0.4402	-0.0668	-0.15	0.6905	0.0953	1		
flowRatio	0.0519	0.6121	-0.1136	-0.1879	-0.0116	0.1064	-0.026	-0.0183	-0.3944	0.0619	-0.3761	1	
absflowratio	0.0388	0.5382	-0.1617	0.0122	-0.0863	0.207	-0.1851	0.2829	-0.2306	0.194	-0.0812	0.6011	1
approach_dominance	0.0519	0.6121	-0.1136	-0.1879	-0.0116	0.1064	-0.026	-0.0183	-0.3944	0.0619	-0.3761	1	0.6011

Other than historical accident data, none of the factors are significantly correlated with either region. As was the case earlier, roundabout inside radius is correlated with outside radius. Low and medium density seem to be moderately correlated with a few factors. This is probably an issue related to sampling size. One or both of these factors might have to be ignored. It is particularly interesting to note that historical accident data has nearly no correlation with traffic volumes or flow ratios.

6.3 Overview and Exploratory Analysis

In the following sections, descriptive statistics are prepared for speed and TTC measures.³

6.3.1 Speed

Comparing hourly speed averages (across the entire merging zone) between Québec and Swedish sites, a drastic reduction in mean speed of about 10 km/h can be seen across all sites (merging zones), even after controlling for traffic volumes at different times of the day as demonstrated in Figure 6.5, and this despite identical posted speed limits (50 km/h).

Similar conclusions can be drawn by comparing the speed profiles for all four merging zone movements between Québec and Sweden, although the effect is observed to be more pro-

³These statistics are automatically generated through tvaLib or tvaLib's included roundabout HLI module.




Figure 6.5 Comparison of hourly speed averages between Québec and Swedish sites.

nounced at the start of each profile, regardless of movement type. Note that in both cases, the speed profiles of movements (i) and (iii), involving a roundabout destination lane, are consistently lower by about 5 km/h than movements (ii) and (iv), involving an exit lane, especially at the end of each profile.

6.3.2 Yielding Post-Encroachment Time

A comparison of yPET is made between Québec and Swedish merging zones. This is illustrated in Figure 6.7. The distribution of both lead and lag yPET below 5 s, measured at the Québec sites is slightly different in amplitude but has a shape similar to the one observed in Chapter 5. Overall, the yPETs below 5 s measured at the Swedish sites are somewhat less clearly distributed, though the general trend of finding peaks in the 2 s, for lag yPET, and 3 s, (for lead yPET, ranges persists. These measures appear to be normally distributed, or at least distributed following a mixture of normal and some negative exponential distributions, though variance is noticeably larger. Furthermore, the mean lead yPET is appreciably higher for the Swedish sites, at about 3 s. The mean lag yPET for Swedish sites is just slightly higher, at 2 s, than for Québec sites.

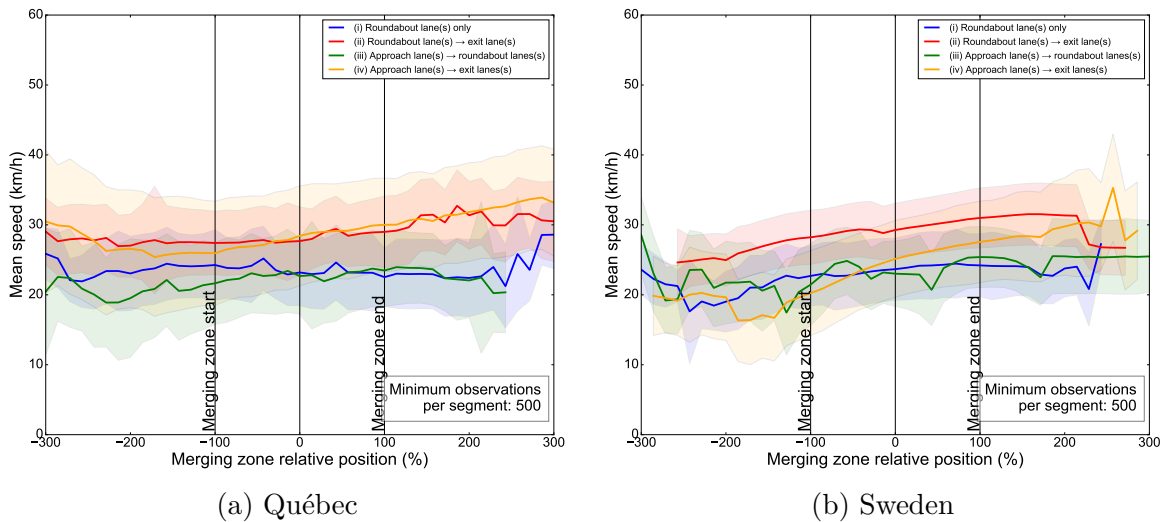


Figure 6.6 Aggregated mean speed and standard deviation profiles by movement types (i) through (iv) for Québec merging zones and Swedish merging zones.

6.3.3 Time-to-Collision

As per the discussion in section 4.6 and methodology applied in Chapter 5, TTC types discretized motion pattern motion prediction and 15th centile aggregation are chosen for safety indicator evaluation, due to their intrinsic benefits in non-linear driving environments, such as in the case of roundabouts, as well as the other benefits of motion pattern motion prediction.

Figure 6.8 compares the cumulative TTC distributions between the Québec sites and the Swedish sites using SCC. Note that a couple of these merging zones have relatively few observations, causing apparent jumps in the cumulative distributions. These Québec sites are included in the results from previous chapters, though they do not appear in distribution plots, as the sample size is too small to appear among sites with considerably more observations. The combined distribution among all observations for all merging zones is plotted with the thick line. Analysis of this data using SCC is arguably inconclusive, but it is shown that a great deal of individual variability can be found between sites. A more thorough statistical regression will be needed to more conclusively interpret the data, especially from individual merging zones.




Figure 6.7 Yielding post-encroachment time distributions for Québec and Swedish merging zones.

6.4 Regression Analysis

In this section, regression analysis is performed following the identical steps outlined in section 5.4 for regression analysis (testing for multicollinearity, homoscedasticity, normality of residuals, model specification, and Wald testing of simple and composite linear hypotheses), although regression for traffic parameters is not performed in this case.

6.4.1 Speed

A stepwise linear regression is performed on mean merging zone speed (measured in km/h) for all road users (`mean_speed`), testing all explainable differences between sites, shown in Tables 6.3 and 6.4. The coefficients, adjusted R^2 , Wald test, and number of observations are provided in Table 6.8.

Note that roundabout radius, flow ratio, land use, urban density, and construction year were not significant in predicting mean speed. Instead, a good model (with an adjusted $R^2 = 0.658$) with only two factors remains:

- A significant reduction in mean speed of 4.5 km/h is observed at the Swedish sites.
- Increases in hourly traffic volume are correlated with reductions in mean speed as well. This is not surprising, given standard traffic flow theory (e.g. Greenshield's Model).

These conclusions are similar to those drawn from the descriptive analysis performed in

section 6.3.1. More importantly, however, is the confirmation that none of the factors that were controlled between sites are shown to have an impact on speed, suggesting that regional, cultural, and road user education effects might be in play instead.

Table 6.8 Linear Regression Models for Mean Speed and Median Lag Yielding Post-Encroachment Time

	Mean Speed		Median Lag yPET	
	Coefficient	P> t	Coefficient	P> t
<code>_cons</code>	35.898	0.000	1.303	0.018
<code>sweden</code>	-4.460	0.007	-	-
<code>outsideRadius_m</code>	-	-	.0798	0.010
<code>inputflowphp1</code>	-0.0240	0.003	-.00336	0.001
Adjusted R^2	0.6582		0.4475	
Wald prob. > F	0.0001		0.0034	
Observations	19		19	

6.4.2 Yielding Post-Encroachment Time

A stepwise linear regression is performed on median lead and lag yPET below 5 s at each site, to test all explainable differences between sites, shown in Tables 6.3 and 6.4. The coefficients, adjusted R^2 , Wald test, and number of observations are provided in Table 6.8.

Region is not found to be significantly correlated with median lead and lag yPET, although the variables `outsideRadius_m` and `inputflowphp1` are associated with lag yPET, having a moderately powerful relationship.

6.4.3 Time-to-Collision

A stepwise linear regression of motion pattern-based serious TTC events (measured in events per hour) is performed, using weighted SEC methodology with a threshold of $\zeta = 1.5$ seconds. No statistically significant model can be found, however. That leaves SCC-based regression as the final analysis to be conducted. As with the TTC regression performed in section 5.4.4, a random effects regression model is justified, using the log of the dependent variable TTC:

$$\ln(TTC_{15ij}) = \alpha + \sum_k \beta_k X_{kij} + u_{ij} + \epsilon_{ij} \quad (6.1)$$

using identical notation as in equation 5.11. A SCC random effects regression of motion pattern-based 15th centile TTC (measured in seconds) is performed between the Québec and

Swedish sites. Results are shown in Table 6.9. This yields a moderately predictive model with a between $R^2 = 0.425$ (accounting for differences between merging zones). This is accounted for by the Swedish variable, as it is associated with increasing TTC by 0.293 s on average, thus leading to hypothetical reductions in collision probability. A minor within-effect is also noted, with fifteen-second traffic exposure associated with an increasing TTC. As evidenced with the angle of incidence parameter, user pairs with a small angle of incidence, i.e. rear-end conflicts, are associated with lower TTC values than with a larger angle of incidence, i.e. side-swipe conflicts. The sign and magnitude of the coefficient associated with this parameter is consistent with the coefficient for angle of incidence observed in Chapter 5. Finally, it is worth noting that the approach dominance and absolute flow ratio parameters are not significantly correlated with TTC as is the case with the Québec roundabout merging zones exclusively. Furthermore, construction year (or elapsed time since roundabout construction) is not correlated significantly either.

Table 6.9 Random Effects Regression Models for Motion Pattern-Based 15th Centile Time-to-Collision (SCC)

	Coefficient	P> t
<code>_cons</code>	0.583	0.000
<code>sweden</code>	0.293	0.029
<code>fifteen_second_exposure</code>	0.01690	0.000
<code>interinst_angle</code>	0.003279	0.000
Within R^2	0.0540	
Between R^2	0.4244	
Overall R^2	0.0204	
Wald prob > F	0.0000	
Observations	23565	
Groups	19	

6.5 Summary

In this chapter, an analysis of roundabout merging zones is performed, as in Chapter 5, but for a different purpose: an international comparison of microscopic road user behaviour and road safety for the purpose of identifying road user behaviour factors that might explain discrepancies in road safety. The same effects on road safety between the two regions are observed consistently across a wide range of measures, reflected in the SSMs of speed, yPET, and TTC, and this after controlling for a number of road geometry, land use, traffic composition, weather and climate conditions, temporal effects, and traffic exposure factors. This leaves a latent, unobserved component of road user behaviour that might be affected by road

culture. Whether this culture is shaped by collective trends in education, enforcement, or design policy, i.e. “culture”, remains to be seen, but it seems clear that road user behaviour is shaped by more than site-based effects.

Furthermore, these observations are consistent with both the macroscopic and microscopic historical accident analysis. Although a more thorough statistical analysis will be necessary to answer this hypothesis more conclusively, the consistency of SSMs observations and interpretations, as well as consistency of these measures with the limited sample of historical accident data, gives validity to the claim that surrogate safety methods can be substituted for historical accident analysis.

One small limitation of this study, however, is that while the sample of Québec roundabouts is regionally representative of most of the province of Québec, the roundabouts sampled in Sweden were limited to the city of Lund. However, this limitation applies only in the situation where the road user behaviour differences within Sweden are greater than those expected between Europe and North America as a whole. In fact, the opposite is assumed.




Figure 6.8 Comparison of TTC distributions between Québec and Sweden. Each individual line is a merging zone while the thick line is the combined distribution of all merging zones of that region.

CHAPTER 7 CONCLUSION

7.1 Summary Conclusions

The work conducted as part of this thesis constitutes a number of practical, theoretical, and empirical contributions to the fields of computer vision applied to traffic analysis, surrogate safety methods, and roundabout design and safety through two roundabout case studies: roundabout design in Québec, and compared between Québec and Sweden. This work is needed because of a lack of systematic definitions of SSMs, lack of systems capable of providing the large quantities of high-resolution data necessary to operate these SSMS, and limited knowledge of roundabout road safety (particularly at the microscopic level) and suitability in North America.

7.1.1 Contributions to Applications of Computer Vision for Traffic Analysis

This work presented a complete, large-scale, microscopic, automated video-based traffic data collection system for general traffic and road safety analysis. This automated video-based traffic data collection system includes an array of small contributions and improvements built upon previous video-based (Bradski, 2000) traffic data collection systems (Jackson et al., 2013). These improvements include

- A continuation of work performed in 2012 (St-Aubin, 2012), and during the course of this project, the design of a VCU (Jackson et al., 2013) specifically for the purpose of collecting the requisite large-quantities of video data affordably.
- A large-volume, low-impact tracking optimization application using MOTA (Bernardin and Stiefelhagen, 2008) to robustly measure tracking performance, and adapted from (Ettahadieh et al., 2015). Site-specific tracking accuracy is measured at up to 94 % (measured on the data used for optimization), while general-purpose tracking parameters are shown to have 80 to 85 % tracking accuracy (as published in Morse et al., 2016) measured through cross-validation (separate optimization and evaluation data).
- A low-volume, large-impact targeted tracking error-correction and filtering quality-control system, and tools for manual video review to remove the errors most problematic to road safety analysis, e.g. erroneously overlapping trajectories. Manual review is an optional component of the automated video-based traffic data collection system.

- Various optimizations to improve performance scaling when processing large volumes of trajectory data from video data, including memory management, a metadata specification used for video sequence indexing, and a host of scene annotation tools.

These are implemented in *tvaLib*, a shared-source library and collection of scripts that provide a variety of tools to automate or semi-automate all steps of the analysis. The library expands the range of functionality of Traffic Intelligence, an open-source, increasingly popular traffic analysis tool in the research community. The most useful components of this software package have already been pack-ported directly into Traffic Intelligence, or are planned for backporting.

7.1.2 Contributions to Surrogate Safety Measures

While the concept of surrogate safety methods (including TCT) is not new, the systematic approach presented as part of this work to calculating, classifying, and aggregating SSMs is innovative and addresses some of the major criticisms present in the field (Williams, 1981; Krussse, 1991; Chin and Quek, 1997). Specifically, original contributions to surrogate safety methods include

- An expanded, complete surrogate safety framework that adapts and applies the existing concepts of the user pair; collision course motion prediction modelling (for new and old models) that forms a unified approach to calculating collision probability from a wide range of motion prediction models; and multiple CPs across multiple interaction instants via SSM aggregation methods. The SEC and SCC analysis approaches are re-contextualized in this framework. Finally, the concept of $P(\textit{evasion})$, and how it relates to SSMs and the surrogate safety framework, is proposed (but not evaluated).
 - The motion prediction methods, constant velocity and normal adaptation, are shown to yield mostly redundant SSM results.
 - Safety indicator aggregation is shown to be sensitive to noise. However, since lack of SSM aggregation causes issues with sampling bias, 15th centile aggregation is proposed as a method of solving some of these issues.
 - Use of SEC methodology has resulted in few useful statistical models in the case study and resulted in inconsistent site rankings for different motion prediction methods. Nevertheless, a modified SEC approach is presented to address the inherent compatibility issues between classic SEC and probabilistic motion prediction models.

- Proposal of a new, flexible, context-aware motion prediction method, the discretized motion pattern motion prediction, with several benefits:
 - motion prediction in non-linear segments of road as are found in roundabouts;
 - incorporation of the empirical prediction of naturalistic road user behaviour common to motion-pattern models;
 - it is simple to implement, and flexible for future addition of context parameters; and
 - it currently implements context parameters representing the origin, speed, and lane of relevant road users.
- SCC-based analysis of TTC measures are consistent with relevant metrics of safety as shown in the international comparison, including pre-collected historical accident data at both a microscopic and macroscopic level, though much more thorough data collection and investigation will be needed in the future in the form of a proper regression analysis using many more sites, preferably with better sources of historical accident data.
- SSMs of speed, yPET, and TTC are shown to be very poorly correlated with one another, suggesting at the very least that these measures represent different dimensions of road safety entirely (e.g. collision probability versus collision severity), or some SSMs may be better suited than others either at predicting dimensions of road safety, or certain dimensions of road safety may have to be subdivided further (i.e. by interaction type, or by magnitude).
- TTC is shown to be generally higher for side-swipe conflict-type interactions than for rear-end conflict-type interactions via the angle of incidence. This is likely explained by the fact that *both* road users have the opportunity to see each other in such situations.
- TTC is shown to decrease when the speed of the slower of the two road users involved in a collision course increases, but is not significantly correlated with an increase in speed of the faster-moving road users. This, however, does not suggest that the differential velocity does not play a factor in TTC.

With this, a recommended TTC-based safety indicator for use in road safety analysis is proposed: discretized motion pattern motion-prediction-based 15th centile TTC aggregation.

7.1.3 Contributions to Roundabout Safety

Use of the merging zone as a unit of analysis is novel, given the microscopic nature of the methodology.

Specific elements of geometry and land use parameters are identified as explaining aspects of road safety. This serves as an example application of the video-based automated traffic data collection system and new the surrogate safety framework. As a result of data collection, 3.7 % of the entire Québec population is sampled. Of these individuals, 8.3 % are found to be involved in a collision course and 1.1 % are found to be involved in at least one serious collision course, according to SEC methodology. Furthermore, the study reaffirms previous suggestions in the literature regarding the interrelationship (see Figure 5.21) between parameters of traffic conditions, road design, and aspects road safety (Ewing and Dumbaugh, 2009; Miranda-Moreno et al., 2011b; Strauss et al., 2013). A number of specific conclusions regarding roundabout design and possibly general road design are drawn as part of the analysis, including, among other conclusions

- Increases in posted speed limits are correlated with increases in speed, though only at a rate of one-tenth the posted speed limit. The range of speed limits tested includes all valid speed limits between 30 km/h and 90 km/h.
- Irregular merging zone configurations are found to be associated with higher speeds and lower (more dangerous) TTC measures.
- While multi-lane roundabouts have been identified as problematic in the literature, few of the models are capable of explaining aspects of road safety from number of lanes directly. Instead, aspects of road safety are modelled as a function of traffic parameters, themselves a function of the number of lanes (by design). This suggests that issues with multi-lane roundabouts in their current form are intractable without reducing traffic volumes.
- Simple traffic circle to roundabout conversions, i.e. without total reconstruction, are associated with lower TTC safety indicators and should be avoided altogether.

An international detailed comparison of road user behaviour and road safety at a microscopic level has never been attempted before. Average operating speeds at Swedish roundabouts was found to be consistently 5 to 10 km/h lower than at similar sites in Québec, and this was after controlling for a number of geometry, land use, and traffic flow parameters, including traffic volumes. These road safety implications are consistent with TTC regression results

which found a statistically significant increase in TTC at Swedish roundabouts, suggesting lower probabilities of collision. Furthermore, these results are consistent with the expected results of surrogate safety theory, as is dictated by national and local-level historical accident data, though not enough to provide a full calibration of the impact of what specific SSMS measure.

7.2 Limitations

7.2.1 Applications of Computer Vision for Traffic Analysis

The feature-based tracker implemented in Traffic Intelligence is not perfect. Several errors persist even after the application of tracking optimization and error correction strategies. Quantifying this error is difficult, since different types of errors exist and their impact can be measured differently, depending on how the data is used and interpreted. General tracking optimization was performed and a MOTA accuracy of between 80 and 85 % was obtained. However, Any remaining tracking errors can still have a disproportionate impact on results, especially since road safety analysis is the very study of outlying road user behaviour. This is the primary purpose of the error-correcting functions: to target and remove tracking issues that have been identified and are known to impact analysis disproportionately. Arguably, the significance of results is therefore somewhat diminished.

However, it is argued that since the error is consistent across all data sources and all calculations are conducted systematically and identically for all samples, then the conclusions of a before-after or cross-sectional study of different sites should not be affected greatly since the error is reproduced consistently between regions. This assumes that the uncontrolled factors (e.g. different region of the world) do not impact tracking results. For example, the feature-based tracker currently underperforms with large vehicles, such as with trucks. If one site disproportionately contains large trucks, then the error is not identical across sites.

Fortunately, for the roundabout case studies presented in Chapters 5 and 6, no tracking error-sensitive dissimilarities (e.g. trucks and buses or poor visibility conditions) are present between samples. For studies where these dissimilarities are necessary, it would be conceivably challenging to apply the proposed methodology in the near future. However, at the current rate of progress of the field of computer vision, tracking errors are expected to be resolved soon. The error correction contributions of this work may even be obsolete by then.

Finally, an important limitation of this video-based traffic data collection system is its dependence on line of sight and on the quality of lighting conditions. However, no other external sensor technology exists that can reliably capture high-fidelity positions of road users con-

tinually. Practical options exist to limit these issues. For example, line of sight issues can be addressed with higher-angled VCU installations or even using aerial drones (as in Puri, 2005; Coifman et al., 2006; Kanistras et al., 2013). Meanwhile, the video-based traffic data collection system is compatible with any type of video data, including optic sensors designed to operate in poor lighting conditions (e.g. thermographic optics, as used in Fu et al., 2016) or even hyperspectral imaging solutions.

7.2.2 Surrogate Safety Measures

The most important limitation of this work is that TTC has yet to be calibrated fully (e.g. for its quantitative effect) as a suitable predictor of the expected number of collisions. For now, use of TTC results is mostly qualitative in nature, and it is substantiated by its broad appeal in the literature, and its intrinsic properties, and by the fact that the application in this work constitutes one of the most robust implementations to date.

The discretized motion pattern motion prediction, as for all types of motion pattern motion prediction, cannot predict collision courses further than the data that is available to learn motion. In practice, this means that motion prediction is bounded by the range of trajectories collected, e.g. the mask. Furthermore, discretized motion pattern motion prediction is subject to outlying road user behaviour sampling issues. Again, this issue is common to all motion pattern motion prediction methods. However, a problem unique to discretized motion pattern motion prediction is the aliasing effect caused by the discretization process, and resulting probability calculation simplifications. Fortunately, all these issues can be addressed with larger data collection.

7.2.3 Roundabout Safety

While use of the merging zone as a unit of analysis is novel, it does impose boundary issues whereby user pairs that otherwise might be negotiating trajectories across multiple merging zones, become isolated. There is especially a small possibility that some of the issues of multi-lane roundabouts are not captured uniquely in the merging zone. Yet, the effect of this problem is limited by the inclusion of significant portions of the approach and exit as part of the studied merging zones (except for yPET—by design). Furthermore, road users that never meet in the same merging zones simultaneously or near-simultaneously are arguably never truly in any significant danger to begin with. Furthermore its use could be problematic for future comparisons between roundabouts and non-roundabout intersections. This issue is to be addressed through careful selection of relevant analysis zones in non-roundabout intersections.

Roundabouts are one of the best environments for road user behaviour comparison between Québec and Sweden, given that roundabout design in North America was imported from Europe and roundabouts are one of the types of road infrastructure with the most similarities between continents. Comparison using other types of road geometry is warranted, in order to study the association of road culture with road user behaviour and road safety.

The sample of Swedish roundabouts is a bit narrow, given that the data collected was concentrated in Lund. An assumption needs to be made that motorists of Lund are representative of the remainder of Sweden for the study to be a true comparison between Québec and Swedish road users. This is somewhat justified, because road users are not stationary—they travel inter-city distances more or less regularly—and some cultural homogeneity can be expected as a product of these individuals interacting with each other at different locations. At the very least, this level of homogeneity may be much more pronounced within a small country than between continents. Furthermore, it should be noted that land use, density, and roundabouts were tightly controlled across samples, leaving room for only human factors to vary between samples. In the context of a comparison between Europe and North America, the same limitation exists with regards to conclusions when Québec and Swedish samples are used as representative samples of driving culture of the respective regions of the globe. Given the homogeneity of road design and culture within each continent, but the striking differences between the two continents (compare, for example, the prevalence and use of stop signs and yield signs between any two communities from North America and Europe), this assumption is arguably very easy to make.

7.3 Future Work

7.3.1 Applications of Computer Vision to Traffic Analysis

The video-based traffic data collection system presented as part of this work was not built to be used in real time, since such systems incur additional operating costs or reduce precision. However, given the rapid pace of development and decreasing costs of computer hardware, and given a relevant use case (e.g. fully automated, intelligent incident detection), the video-based traffic data collection system presented here can be optimized to provide real-time or quasi-real-time traffic analysis for targeted locations. In the meantime, while current measures of accuracy are “adequate,” there is still room for improvement, be it better cameras, better tracking algorithms, better tracking parameters, more situation-specific tracking parameters, or more intelligent error correction schemes.

The video-based traffic data collection system developed as part of this work integrates a

fully fledged data management framework. This was built partially on the current need to separate individual video recordings from analysis zones. However, it was also done for the future need to use multiple cameras simultaneously to perform traffic and safety analysis on larger analysis zones. One missing piece that remains before such a system is complete is the inclusion of a tool to reconnect trajectories from distinct cameras together, to be processed as a single data sequence. This is the general problem of matching segmented trajectories. Some work has been done in this field, (e.g. Sankararaman et al., 2013) especially with regards to *tracklets*—tracking based on short tracks from weak trackers that yield partial trajectories with matching algorithms such as DTW (e.g. Raptis and Soatto, 2010)—although this has yet to be formally implemented for the trajectories obtained from this video-based traffic data collection system. In addition to trajectory connecting using distinct video sequences, the same methods could be applied to improve error-correction functions for tracking issues related to occlusion and stopped road users. Such a system will likely be implemented in a future iteration of tvaLib.

7.3.2 Surrogate Safety Measures

The benefits of SSMs for road safety analysis are potentially radical, but much more work remains to be done in the field before the methodology is robust enough to be used by practitioners. The chief need in this area is empirical calibration of most SSMs to yield expected numbers of collision. The primary hurdle in this regard is the same as with any historical accident analysis: collection of reliable and purpose-built historical accident data. Fortunately, this only needs to be done once or twice per model. This is in the research mandate of the ongoing InDeV research project involving vulnerable road users (Laureshyn et al., 2015).

Motion Prediction

Motion prediction models can be easily interchanged in the surrogate safety method framework presented in Chapter 4 as implemented in tvaLib. Many more motion prediction models were proposed as feasible methods of motion prediction during the planning stages of the project, including

- The complexity of the discretized motion pattern method presented in section 4.3.2 is limited only by the size of the available training data. With a larger data set, additional parameters could be incorporated, including number of road users in proximity of the vehicle (which would model constrained movement options for road users), visual

obstruction from moving objects, sun blindness, etc.

- A probabilistic approach to motion prediction modelling could guide objects to predicted destinations using lane alignments or prototypes as “rails” (similar to Mohamed and Saunier, 2015), travelling at constant speed (or predicted speed). This offers the intrinsic advantages of motion patterns without the large sampling issues, however this type of model would make more assumptions than motion pattern models.
- A trajectory deviation model could upgrade constant-velocity-based models by accounting for tendencies in lane drift by measuring the rate of change of the lane offset. However, this model would require superior alignments, as it would be sensitive to the quality of the alignments.

Furthermore, motion prediction methods could be expanded to calculate collisions not only between pairs of road users, but also between road users and physical obstacles and barriers in the environment. However, this type of modelling cannot be handled by motion pattern-based motion prediction (as there are no patterns of motion through static objects to begin with). Constant-velocity-based models might excel at this task, as this type of collision might usually be the result of a loss of control. In this case, by the strictest definition, user pair collisions are no longer being modelled, instead, “single-user collisions” are being modelled (separately).

Discretized motion pattern motion prediction and similar empirical motion prediction models will need more fine-grained (more continuous discretization) analysis and calibration, and possibly a sensitivity analysis. It is possible to expand this model to include more factors and thus make fewer assumptions about predicted collision courses, however, this will require more data and additional investigation.

Roundabouts

Although a significant number of studies have been performed on roundabout capacity and operational performance, the automated, large-scale, microscopic traffic data collection system presented as part of this work to investigate further nuances in these design criteria presents an opportunity to revisit these studies. For example, a large number of capacity equations are available that take as input traffic volume for the approach and conflicting flow, and number of lanes (Rodegerdts et al., 2010, e.g.), but do not necessarily consider arrival groupings or effects that may be derived from movement types within the merging zone, or examine how the use of large single lanes by two road users simultaneously may have

a significant impact. In addition, a more in-depth study of multi-lane roundabouts between merging zone is warranted.

Analysis will continue in the short term. Additional data was collected for clear summer and winter storm conditions at one roundabout (6 months apart), and at another roundabout, data was collected immediately prior to and six months after conversion from a traffic-light-controlled intersection. Future work will attempt to (i) quantify the effects of weather on road user behaviour, especially in roundabouts (in a way similar to an analysis performed in Fu et al., 2015), and (ii) complement the cross-sectional analysis presented as part of this work with a classic before/after comparison of roundabout conversion, at a microscopic scale. The comparison of merging zones with traditional areas of a non-roundabout intersection will have to be addressed, but ultimately it does not constitute a serious limitation to the methodology. Finally, analysis of SSM between types of user pair classification (rear-end, side-swipe conflict, etc.) is also planned for future work.

7.4 Final Perspective

The large-scale, automated, microscopic traffic data collection system developed in this research provides the necessary data to perform the systematic measurement required for a robust, unified surrogate safety method framework, while addressing issues with variable SSM definition and measurement, and SSM transferability between studies. With this framework, more proactive road safety analysis may be conducted, though, while results are consistent with historical accident data, more thorough validation of individual SSM remains necessary. Additional uses for this large volume of microscopic data could include calibration of microsimulation and car-following models, automated traffic incident or infraction detection, and a variety of case-specific research uses.

A number of different traffic parameters, land-use, and geometry factors are successfully identified as contributing to roundabout safety in a number of different ways. Furthermore, latent cultural road user behaviour factors are shown to exist between regions of the world. This would suggest that transferring road designs directly between regions of the world (such as was the case with the transplantation of roundabout design from Europe into North America) may be less desirable than adapting designs to accommodate local road user behaviour. What is less clear is to what extent road user behaviour can be adapted solely by changing or implementing new road design.

BIBLIOGRAPHY

AASHTO, “A policy on geometric design of highways and streets,” American Association of State Highway and Transportation Officials, Washington, DC, Tech. Rep., 2004.

K. I. Ahmed, “Modeling drivers’ acceleration and lane changing behavior,” Ph.D. dissertation, Department of Civil and Environmental Engineering, Cambridge, MA, 1999.

M. Al-Ghandour, “Roundabout Slip Lanes: Performance and Safety Analysis,” Ph.D. dissertation, Raleigh, 2011.

B. L. Allen, B. T. Shin, and P. J. Cooper, “Analysis of traffic conflicts and collisions,” *Transportation Research Record: Journal of the Transportation Research Board*, vol. 667, pp. 67–74, 1978.

S. Almqvist and L. Ekman, “The Swedish Traffic Conflict Technique Observers’ Manual,” Lund University, Lund, Sweden, Tech. Rep., 2001.

M. Althoff, O. Stursberg, and M. Buss, “Stochastic reachable sets of interacting traffic participants,” in *Intelligent Vehicles Symposium, 2008 IEEE*. IEEE, 2008, pp. 1086–1092.

American Association of State Highway and Transportation Officials, *Highway Safety Manual Three-volume Set 2010*. American Association of State Highway and Transportation Officials, 2010.

F. Amundsen and C. Hydén, “Proceedings of the first workshop on traffic conflicts,” in *Institute of Transport Economics*, Oslo, Norway, 1977, p. 76, the famous “unchanging” TTC quote.

J. Archer, “Methods for the assessment and prediction of traffic safety at urban intersections and their application in micro-simulation modelling,” Ph.D. dissertation, Sockholm, Sweden, 2004.

J. Autey, T. Sayed, and M. H. Zaki, “Safety evaluation of right-turn smart channels using automated traffic conflict analysis,” *Accident Analysis & Prevention*, vol. 45, pp. 120–130, Mar 2012. [Online]. Available: <http://dx.doi.org/10.1016/j.aap.2011.11.015>

S. P. Baker, B. O’Neill, W. Haddon, and W. B. Long, “The Injury Severity Score,” *The Journal of Trauma: Injury, Infection, and Critical Care*, vol. 14, no. 3, pp. 187–196, mar 1974. [Online]. Available: <http://dx.doi.org/10.1097/00005373-197403000-00001>

M. Bennewitz, W. Burgard, G. Cielniak, and S. Thrun, “Learning motion patterns of people for compliant robot motion,” *International Journal of Robotics Research*, vol. 24, pp. 31–48, 2005.

M. Bennewitz, W. Burgard, G. Cielniak, and S. Thrun, “Learning motion patterns of people for compliant robot motion,” *The International Journal of Robotics Research*, vol. 24, no. 1, pp. 31–48, 2005.

K. Bernardin and R. Stiefelhagen, “Evaluating multiple object tracking performance: the CLEAR MOT metrics,” *Journal on Image and Video Processing*, vol. 2008, p. 1, 2008.

R. Bhatia and M. Wier, ““Safety in Numbers” re-examined: Can we make valid or practical inferences from available evidence?” *Accident Analysis & Prevention*, vol. 43, no. 1, pp. 235–240, jan 2011, safety in numbers critic. [Online]. Available: <http://dx.doi.org/10.1016/j.aap.2010.08.015>

C. E. Bonferroni, “Teoria statistica delle classi e calcolo delle probabilita.” (Pubbl. d. R. Ist. Super. di Sci. Econom. e Commerciali di Firenze. 8) Firenze: Libr. Internaz. Seeber. 62 S. (1936)., 1936.

M. Brackstone and M. McDonald, “Car-following: a historical review,” *Transportation Research Part F: Traffic Psychology and Behaviour*, vol. 2, no. 4, pp. 181–196, dec 1999. [Online]. Available: [http://dx.doi.org/10.1016/S1369-8478\(00\)00005-X](http://dx.doi.org/10.1016/S1369-8478(00)00005-X)

G. Bradski, “OpenCV Library,” *Dr. Dobb’s Journal of Software Tools*, 2000.

S. Brahmabhatt, *Practical OpenCV*. Apress, 2013. [Online]. Available: <http://dx.doi.org/10.1007/978-1-4302-6080-6>

M. D. Breitenstein, F. Reichlin, B. Leibe, E. Koller-Meier, and L. Van Gool, “Online multiperson tracking-by-detection from a single, uncalibrated camera,” *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 33, no. 9, pp. 1820–1833, 2011.

G. R. Brown, “Traffic conflicts for road user safety studies,” *Canadian Journal of Civil Engineering*, vol. 21, no. 1, pp. 1–15, feb 1994. [Online]. Available: <http://dx.doi.org/10.1139/194-001>

U. Brüde and J. Larsson, “Models for predicting accidents at junctions where pedestrians and cyclists are involved. How well do they fit?” *Accident Analysis & Prevention*, vol. 25, no. 5, pp. 499–509, oct 1993. [Online]. Available: [http://dx.doi.org/10.1016/0001-4575\(93\)90001-d](http://dx.doi.org/10.1016/0001-4575(93)90001-d)

- N. Buch, S. Velastin, and J. Orwell, "A Review of Computer Vision Techniques for the Analysis of Urban Traffic," *Intelligent Transportation Systems, IEEE Transactions on*, vol. 12, no. 3, pp. 920–939, Sept 2011. [Online]. Available: <http://dx.doi.org/10.1109/TITS.2011.2119372>
- J. Bull and B. Roberts, "Road accident statistics—A comparison of police and hospital information," *Accident Analysis & Prevention*, vol. 5, no. 1, pp. 45–53, apr 1973. [Online]. Available: [http://dx.doi.org/10.1016/0001-4575\(73\)90004-3](http://dx.doi.org/10.1016/0001-4575(73)90004-3)
- S. Burns, L. Miranda-Moreno, N. Saunier, and K. Ismail, "Crash Severity Analysis at Roundabouts: Case Study in Quebec, Canada," in *Transportation Research Board 92nd Annual Meeting*. Washington, D.C.: National Academy of Sciences, 2013, p. 14.
- Y. Chen, B. Persaud, E. Sacchi, and M. Bassani, "Investigation of models for relating roundabout safety to predicted speed," *Accident Analysis & Prevention*, vol. 50, pp. 196–203, jan 2013. [Online]. Available: <http://dx.doi.org/10.1016/j.aap.2012.04.011>
- H.-C. Chin and S.-T. Quek, "Measurement of traffic conflicts," *Safety Science*, vol. 26, no. 3, pp. 169–185, Aug 1997. [Online]. Available: [http://dx.doi.org/10.1016/s0925-7535\(97\)00041-6](http://dx.doi.org/10.1016/s0925-7535(97)00041-6)
- B. Coifman, M. McCord, R. Mishalani, M. Iswalt, and Y. Ji, "Roadway traffic monitoring from an unmanned aerial vehicle," in *IEE Proceedings-Intelligent Transport Systems*, vol. 153, no. 1. IET, 2006, pp. 11–20.
- B. Coifman, D. Beymer, P. McLauchlan, and J. Malik, "A real-time computer vision system for vehicle tracking and traffic surveillance," *Transportation Research Part C: Emerging Technologies*, vol. 6, no. 4, pp. 271–288, 1998.
- B. Cumming, "High rate of crashes at roundabouts involving cyclists can be reduced with careful attention to conflict paths," *Injury prevention*, vol. 18, no. 1, p. A24, 2012.
- S. Daniels, T. Brijs, E. Nuyts, and G. Wets, "Explaining variation in safety performance of roundabouts," *Accident Analysis & Prevention*, vol. 42, no. 2, pp. 393–402, mar 2010. [Online]. Available: <http://dx.doi.org/10.1016/j.aap.2009.08.019>
- G. A. Davis, J. Hourdos, and H. Xiong, "Outline of causal theory of traffic conflicts and collisions," in *Transportation Research Board 87th Annual Meeting*, no. 08-2431, 2008.
- A. Eidehall and L. Petersson, "Statistical Threat Assessment for General Road Scenes Using Monte Carlo Sampling," *IEEE Trans. Intell. Transport. Syst.*, vol. 9, no. 1, pp. 137–147, mar 2008. [Online]. Available: <http://dx.doi.org/10.1109/TITS.2007.909241>

- R. Elvik, "The Power Model of the relationship between speed and road safety: update and new analyses," The Institute of Transport Economics TØI, Oslo, Norway, Tech. Rep. 1034/2009, 2009.
- R. Elvik, P. Christensen, and A. Amundsen, "Speed and Road Accidents: An evaluation of the Power Model," The Institute of Transport Economics TØI, Oslo, Norway, Tech. Rep. 740/2004, 2004.
- R. Elvik, A. Erke, and P. Christensen, "Elementary Units of Exposure," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2103, pp. 25–31, dec 2009. [Online]. Available: <http://dx.doi.org/10.3141/2103-04>
- R. Ervin, C. MacAdam, J. Walker, S. Bogard, M. Hagan, A. Vayda, and E. Anderson, "System for Assessment of the Vehicle Motion Environment (SAVME)," Tech. Rep. UMTRI-00-21, 2000.
- D. Ettihadieh, B. Farooq, and N. Saunier, "Systematic Parameter Optimization and Application of Automated Tracking in Pedestrian-Dominant Situations," in *trb*, 2015, 15-2400.
- R. Ewing and E. Dumbaugh, "The Built Environment and Traffic Safety: A Review of Empirical Evidence," *Journal of Planning Literature*, vol. 23, no. 4, pp. 347–367, may 2009. [Online]. Available: <http://dx.doi.org/10.1177/0885412209335553>
- B. Fildes and S. Lee, "The Speed Review: Road Environment, Behaviour, Speed limits, Enforcement and Crashes," Roads and Traffic Authority of New South Wales, Victoria, Australia, Tech. Rep., 1993.
- T. Fu, S. Zangenehpour, P. St-Aubin, L. Fu, and L. F. Miranda-Moreno, "Using microscopic video data measures for driver behavior analysis during adverse winter weather: Opportunities and challenges," *Journal of Modern Transportation*, vol. 23, no. 2, pp. 81–92, may 2015. [Online]. Available: <http://dx.doi.org/10.1007/s40534-015-0073-3>
- T. Fu, J. Stipanovic, S. Zangenehpour, L. Miranda-Moreno, and N. Saunier, "A Comparison of Regular and Thermal Cameras for Traffic Data Collection under Varying Lighting and Temperature Conditions in Multimodal Environments," in *Transportation Research Board (TRB) 95th Annual Meeting*, Transportation Research Board, Ed. Washington, D.C.: National Academy of Sciences, 2016, p. 18.
- D. Gettman and L. Head, "Surrogate Safety Measures From Traffic Simulation Models," p. 118, 2003.

D. Gettman, L. Pu, T. Sayed, and S. Shelby, "Surrogate Safety Assessment Model and Validation: Final Report," Federal Highway Administration, Tech. Rep., june 2008.

T. Gordon, Z. Bareket, L. Kostyniuk, M. Barnes, M. Hagan, Z. Kim, D. Cody, A. Skabardonis, and A. Vayda, "Site-Based Video System Design and Development," Strategic Highway Research Program (SHRP2), Tech. Rep. S2-S09-RW-1, 2012.

U. Grayson, "Traffic Conflicts in Britain: The Past and The Future," in *International Calibration Study of Traffic Conflict Techniques*. Springer, 1984, pp. 25–33.

M. Green, "How Long Does It Take to Stop? Methodological Analysis of Driver Perception-Brake Times," *Transportation Human Factors*, vol. 2, no. 3, pp. 195–216, Sep 2000. [Online]. Available: http://dx.doi.org/10.1207/STHF0203_1

F. Gross, C. Lyon, B. Persaud, and R. Srinivasan, "Safety effectiveness of converting signalized intersections to roundabouts," *Accident Analysis & Prevention*, vol. 50, pp. 234–241, jan 2013. [Online]. Available: <http://dx.doi.org/10.1016/j.aap.2012.04.012>

G. Guido, F. Saccomanno, A. Vitale, V. Astarita, and D. Festa, "Comparing Safety Performance Measures Obtained from Video Capture Data," *Journal of Transportation Engineering*, vol. 137, no. 7, pp. 481–491, jul 2011. [Online]. Available: [http://dx.doi.org/10.1061/\(ASCE\)TE.1943-5436.0000230](http://dx.doi.org/10.1061/(ASCE)TE.1943-5436.0000230)

A. S. Hakkert, L. Braimaister, and I. van Schagen, "The uses of exposure and risk in road safety studies," vol. 2002, no. 12, 2002.

A. Hakkert, "Review of Traffic Conflict Technique Applications in Israel," in *International Calibration Study of Traffic Conflict Techniques*. Springer, 1984, pp. 7–16.

E. Hauer, "Traffic conflicts and exposure," *Accident Analysis & Prevention*, vol. 14, no. 5, pp. 359–364, oct 1982. [Online]. Available: [http://dx.doi.org/10.1016/0001-4575\(82\)90014-8](http://dx.doi.org/10.1016/0001-4575(82)90014-8)

E. Hauer, "Speed and Safety," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2103, no. -1, pp. 10–17, Sep 2009. [Online]. Available: <http://dx.doi.org/10.3141/2103-02>

J. Hayward, "Near misses as a measure of safety at urban intersections," Ph.D. dissertation, 1971, first use of TTC.

W. Hu, X. Xiao, D. Xie, T. Tan, and S. Maybank, "Traffic Accident Prediction Using 3-D Model-Based Vehicle Tracking," *IEEE Trans. Veh. Technol.*, vol. 53, no. 3, pp. 677–694, may 2004. [Online]. Available: <http://dx.doi.org/10.1109/TVT.2004.825772>

C. Hupfer, “Deceleration to Safety Time (DST) - a Useful Figure to Evaluate Traffic Safety.” Lund, Sweden: Department of Traffic Planning and Engineering, 1997, first use of DST.

C. Hydén, “The Development of a Method for Traffic Safety Evaluation: The Swedish Traffic Conflicts Technique,” Ph.D. dissertation, Department of Traffic Planning and Engineering, 1987, tCT thresholds, conversion rates, and reliability tests.

C. Hydén and A. Várhelyi, “The effects on safety time consumption and environment of large scale use of roundabouts in an urban area: A case study,” *Accident Analysis & Prevention*, vol. 32, no. 1, pp. 11–23, jan 2000. [Online]. Available: [http://dx.doi.org/10.1016/s0001-4575\(99\)00044-5](http://dx.doi.org/10.1016/s0001-4575(99)00044-5)

Institute for Road Safety Research, “SWOV Fact sheet - The relation between speed and crashes,” Leidschendam, the Netherlands, Tech. Rep., 2012.

J. Irvenå and S. Randahl, “Analysis of gap acceptance in a saturated two-lane roundabout and implementation of critical gaps in VISSIM,” Master’s thesis, Department of Technology and Society, Lund, Sweden, 2010.

H. Isebrands, “Quantifying safety and speed data for rural roundabouts with high-speed approaches,” Ph.D. dissertation, Ames, Iowa, 2011.

K. Ismail, T. Sayed, N. Saunier, and C. Lim, “Automated Analysis of Pedestrian-Vehicle Conflicts using Video Data,” *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2140, pp. 44–54, 2009. [Online]. Available: <http://dx.doi.org/10.3141/2140-05>

K. Ismail, T. Sayed, and N. Saunier, “Automated Analysis of Pedestrian-Vehicle Conflicts: Context For Before-and-after Studies,” *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2198, pp. 52–64, 2010.

ISO 31000, *Risk management - Principles and guidelines*, 2009, no. ISO 31000:2009.

S. Jackson, L. F. Miranda-Moreno, P. St-Aubin, and N. Saunier, “Flexible, Mobile Video Camera System and Open Source Video Analysis Software for Road Safety and Behavioral Analysis,” *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2365, no. -1, pp. 90–98, Dec 2013. [Online]. Available: <http://dx.doi.org/10.3141/2365-12>

P. L. Jacobsen, “Safety in numbers: more walkers and bicyclists, safer walking and bicycling,” *Injury Prevention*, vol. 9, no. 3, pp. 205–209, sep 2003, safety in numbers. [Online]. Available: <http://dx.doi.org/10.1136/ip.9.3.205>

- G. Jacquemart, "Synthesis of Highway Practice 264: Modern Roundabout Practice in the United States," National Cooperative Highway Research Program, Washington, D.C., Tech. Rep., 1998.
- S. U. Jensen, "Safety Effects of Converting Intersections to Roundabouts," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2389, pp. 22–29, dec 2013. [Online]. Available: <http://dx.doi.org/10.3141/2389-03>
- S. U. Jensen, "Safety Effects of Height of Central Islands, Sight Distances, Markings and Signage at Single-lane Roundabouts," in *Transportation Research Board (TRB) 93th Annual Meeting*, Transportation Research Board, Ed. Washington, D.C.: National Academy of Sciences, 2014, p. 16.
- J.-P. Jodoin, G.-A. Bilodeau, and N. Saunier, "Tracking All Road Users at Multimodal Urban Traffic Intersections," *IEEE Trans. Intell. Transport. Syst.*, pp. 1–11, 2016. [Online]. Available: <http://dx.doi.org/10.1109/TITS.2016.2545245>
- E. Jones, T. Oliphant, P. Peterson *et al.*, "SciPy: Open source scientific tools for Python," 2001–, [Online; accessed 2015-12-15]. [Online]. Available: <http://www.scipy.org/>
- K. Kanistras, G. Martins, M. J. Rutherford, and K. P. Valavanis, "A survey of unmanned aerial vehicles (UAVs) for traffic monitoring," in *Unmanned Aircraft Systems (ICUAS), 2013 International Conference on*, May 2013, pp. 221–234. [Online]. Available: <http://dx.doi.org/10.1109/ICUAS.2013.6564694>
- Z. Kim, G. Gomes, R. Hranac, and A. Skabardonis, "A Machine Vision System for Generating Vehicle Trajectories over Extended Freeway Segments," in *12th World Congress on Intelligent Transportation Systems*, San Francisco, CA, nov 2005, nGSIM.
- J. Kraay, A. van der Horst, and S. Oppe, "Manual conflict observation technique DOCTOR-Dutch Objective Conflict Technique for Operation and Research," Voorburg, The Netherlands, p. 86, 2013.
- H. W. Kruijse, "The subjective evaluation of traffic conflicts based on an internal concept of dangerousness," *Accident Analysis & Prevention*, vol. 23, no. 1, pp. 53–65, Feb 1991. [Online]. Available: [http://dx.doi.org/10.1016/0001-4575\(91\)90035-4](http://dx.doi.org/10.1016/0001-4575(91)90035-4)
- Y. Kuang, X. Qu, J. Weng, and A. Etemad-Shahidi, "How Does the Driver's Perception Reaction Time Affect the Performances of Crash Surrogate Measures?" *PLoS One*, vol. 10, p. e0138617, 2015.

C. Laugier, I. E. Paromtchik, M. Perrollaz, M. Y. Yong, J. Yoder, C. Tay, K. Mekhnacha, and A. Negre, “Probabilistic Analysis of Dynamic Scenes and Collision Risks Assessment to Improve Driving Safety,” *IEEE Intelligent Transportation Systems Magazine*, vol. 3, no. 4, pp. 4–19, 2011. [Online]. Available: <http://dx.doi.org/10.1109/MITS.2011.942779>

A. Laureshyn, H. Ardö, Å. Svensson, and T. Jonsson, “Application of automated video analysis for behavioural studies: concept and experience,” *IET Intelligent Transport Systems*, vol. 3, no. 3, pp. 345–357, September 2009. [Online]. Available: <http://dx.doi.org/10.1049/iet-its.2008.0077>

A. Laureshyn, “Application of automated video analysis to road user behaviour,” Ph.D. dissertation, Department of Technology and Society, Lund, Sweden, 2010.

A. Laureshyn, “T-Analyst software,” <http://www.tft.lth.se/video/cooperation/software/>, 2013.

A. Laureshyn, Å. Svensson, and C. Hydén, “Evaluation of traffic safety, based on micro-level behavioural data: Theoretical framework and first implementation,” *Accident Analysis & Prevention*, vol. 42, no. 6, pp. 1637–1646, Nov 2010. [Online]. Available: <http://dx.doi.org/10.1016/j.aap.2010.03.021>

A. Laureshyn, A. Varhelyi, and Å. Svensson, “InDeV: In-Depth understanding of accident causation for Vulnerable road users - Project Plan,” Lund University, Sweden, Lund, Sweden, Tech. Rep., 2015. [Online]. Available: http://www.indev-project.eu/InDeV/EN/Home/home_node.html

L. Leden, P. Gårder, and U. Pulkkinen, “An expert judgment model applied to estimating the safety effect of a bicycle facility,” *Accident Analysis & Prevention*, vol. 32, no. 4, pp. 589–599, jul 2000. [Online]. Available: [http://dx.doi.org/10.1016/s0001-4575\(99\)00090-1](http://dx.doi.org/10.1016/s0001-4575(99)00090-1)

Y. Lou, C. Zhang, Y. Zheng, X. Xie, W. Wang, and Y. Huang, “Map-matching for low-sampling-rate GPS trajectories,” in *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems - GIS '09*. Association for Computing Machinery (ACM), 2009. [Online]. Available: <http://dx.doi.org/10.1145/1653771.1653820>

J. B. MacQueen, “Some Methods for Classification and Analysis of Multivariate Observations,” in *5-th Berkeley Symposium on Mathematical Statistics and Probability*. University of California Press, 1967, pp. 1:281–297.

D. L. Massie and K. L. Campbell, “Analysis of Accident Rates by Age, Gender, and Time of Day Based on the 1990 Nationwide Personal Transportation Survey-Final report,” University of Michigan Transportation Research Institute, Tech. Rep., 1993.

B. Maurin, O. Masoud, and N. Papanikolopoulos, “Tracking all traffic: computer vision algorithms for monitoring vehicles, individuals, and crowds,” *Robotics Automation Magazine, IEEE*, vol. 12, no. 1, pp. 29–36, March 2005. [Online]. Available: <http://dx.doi.org/10.1109/MRA.2005.1411416>

T. H. Maze, “A probabilistic model of gap acceptance behavior,” *Transportation research record*, vol. 795, pp. 8–13, 1981.

D. R. Mazur, *Combinatorics: A Guided Tour (MAA Textbooks)*. Mathematical Association of America, 2009.

P. Michalopoulos, “Vehicle detection video through image processing: the Autoscope system,” *IEEE Trans. Veh. Technol.*, vol. 40, no. 1, pp. 21–29, feb 1991. [Online]. Available: <http://dx.doi.org/10.1109/25.69968>

S. Midenet, N. Saunier, and F. Boillot, “Exposure to lateral collision in signalized intersections with protected left turn under different traffic control strategies,” *Accident Analysis & Prevention*, vol. 43, no. 6, pp. 1968–1978, nov 2011. [Online]. Available: <http://dx.doi.org/10.1016/j.aap.2011.05.015>

J. M. Milla, F. Barrero, M. Vargas, and S. L. Toral, *Computer Vision Techniques for Background Modeling in Urban Traffic Monitoring*. INTECH Open Access Publisher, 2010.

Ministère des Transports du Québec, “Le carrefour giratoire: un mode de gestion différent,” Direction du soutien à l’exploitation des infrastructures, Montréal, QC, Tech. Rep., 2002.

L. F. Miranda-Moreno, P. Morency, and A. M. El-Geneidy, “The link between built environment, pedestrian activity and pedestrian–vehicle collision occurrence at signalized intersections,” *Accident Analysis & Prevention*, vol. 43, no. 5, pp. 1624–1634, sep 2011. [Online]. Available: <http://dx.doi.org/10.1016/j.aap.2011.02.005>

L. F. Miranda-Moreno, J. Strauss, and P. Morency, “Disaggregate Exposure Measures and Injury Frequency Models of Cyclist Safety at Signalized Intersections,” *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2236, pp. 74–82, dec 2011. [Online]. Available: <http://dx.doi.org/10.3141/2236-09>

J. Moeckli and J. D. Lee, “The making of driving cultures,” *Improving Traffic Safety Culture in the United States*, vol. 38, no. 2, pp. 185–192, 2007.

M. G. Mohamed, “Automatic Behavior Analysis and Understanding of Collision Processes Using Video Sensors,” Ph.D. dissertation, Département des génies civil, géologique et des mines, Montréal, Canada, 2015. [Online]. Available: <http://publications.polymtl.ca/1784/>

M. G. Mohamed and N. Saunier, “Behavior Analysis Using a Multilevel Motion Pattern Learning Framework,” *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2528, pp. 116–127, sep 2015. [Online]. Available: <http://dx.doi.org/10.3141/2528-13>

M. G. Mohamed and N. Saunier, “Motion Prediction Methods for Surrogate Safety Analysis,” *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2386, pp. 168–178, Dec 2013. [Online]. Available: <http://dx.doi.org/10.3141/2386-19>

B. Morris and M. Trivedi, “A Survey of Vision-Based Trajectory Learning and Analysis for Surveillance,” *Circuits and Systems for Video Technology, IEEE Transactions on*, vol. 18, no. 8, pp. 1114–1127, August 2008. [Online]. Available: <http://dx.doi.org/10.1109/TCSVT.2008.927109>

P. Morse, P. St-Aubin, L. F. Miranda-Moreno, and N. Saunier, “Transferability Study of Video Tracking Optimization for Traffic Data Collection and Analysis,” in *Transportation Research Board (TRB) 95th Annual Meeting*, Transportation Research Board, Ed. Washington, D.C.: National Academy of Sciences, Jan 2016.

National Safety Council, “Estimating the Costs of Unintentional Injuries 2011,” Itasca, Illinois, Tech. Rep., 2013. [Online]. Available: <http://www.nsc.org/Membership%20Site%20Document%20Library/NSCLibrary/Estimating-the-Cost-of-Unintentional-Injuries,-2011.pdf>

NHTSA, “Traffic Safety Facts 1996: Speeding,” U.S. Department of Transportation, Tech. Rep., 1997.

OECD, *ITF Transport Outlook 2015*. Organisation for Economic Co-Operation and Development, jan 2015. [Online]. Available: <http://dx.doi.org/10.1787/9789282107782-en>

OECD and ITF, *Road Safety Annual Report 2015 - Summary*, Paris, France, 2015.

M. Parker and C. Zegeer, “Traffic Conflict Techniques for Safety and Operations - Observers’ Manual,” p. 36, 1989.

- E. Pasanen, "Driving speeds and pedestrian safety: a mathematical model," Helsinki University of Technology, Tech. Rep. Publication 77, 1992.
- L. Pellecuer and M. St-Jacques, "Dernières avancées sur les carrefours giratoires," *Canadian Journal of Civil Engineering*, vol. 35, no. 5, pp. 542–553, may 2008. [Online]. Available: <http://dx.doi.org/10.1139/L07-125>
- M. Perdomo, "Visual Approaches to Understanding Pedestrian Safety in Roundabouts," Master's thesis, Geography, Planning and Environmental Studies, Montreal, QC, 2015.
- M. Perdomo, A. Rezaei, Z. Patterson, N. Saunier, and L. F. Miranda-Moreno, "Pedestrian preferences with respect to roundabouts—A video-based stated preference survey," *Accident Analysis & Prevention*, vol. 70, pp. 84–91, sep 2014. [Online]. Available: <http://dx.doi.org/10.1016/j.aap.2014.03.010>
- S. Perkins and J. Harris, "Traffic conflicts characteristics: Accident potential at intersections," *Highway Research Record*, vol. 225, pp. 35–43, 1968.
- B. Persaud, R. Retting, P. Garder, and D. Lord, "Safety Effect of Roundabout Conversions in the United States: Empirical Bayes Observational Before-After Study," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 1751, pp. 1–8, jan 2001. [Online]. Available: <http://dx.doi.org/10.3141/1751-01>
- PIARC Technical Committee on Road Safety, *Road Safety Manual*. Route2Market, 2004. [Online]. Available: <http://dx.doi.org/10.1007/978-1-4302-6080-6>
- P. A. Pisano, L. C. Goodwin, and M. A. Rossetti, "US Highway Crashes in Adverse Road Weather Conditions," in *4th Conference on International Interactive Information and Processing Systems for Meteorology, Oceanography and Hydrology*, New Orleans, LA, 2008.
- V. Punzo, M. T. Borzacchiello, and B. Ciuffo, "Estimation of vehicle trajectories from observed discrete positions and Next-Generation Simulation Program (NGSIM) data," in *Transportation Research Board (TRB) 88th Annual Meeting*, Transportation Research Board, Ed. Washington, D.C.: National Academy of Sciences, 2009, p. 217.
- A. Puri, "A Survey of Unmanned Aerial Vehicles (UAV) for Traffic Surveillance," University of South Florida, Tech. Rep., 2005.
- X. Qin, J. N. Ivan, and N. Ravishanker, "Selecting exposure measures in crash rate prediction for two-lane highway segments," *Accident Analysis & Prevention*, vol. 36, no. 2, pp. 183–191, 2004.

H. Rakha, B. Hellinga, M. Van Aerde, and W. Perez, “Systematic verification, validation and calibration of traffic simulation models,” in *Transportation Research Board (TRB) 75th Annual Meeting*, Transportation Research Board, Ed. Washington, D.C.: National Academy of Sciences, 1996, p. 14.

M. Raptis and S. Soatto, “Tracklet descriptors for action modeling and video analysis,” in *Computer Vision—ECCV 2010*. Springer, 2010, pp. 577–590.

R. Retting, S. Kyrychenko, and A. McCartt, “Long-Term Trends in Public Opinion Following Construction of Roundabouts,” *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2019, pp. 219–224, dec 2007. [Online]. Available: <http://dx.doi.org/10.3141/2019-26>

E. Rice, “Roundabouts: Technical Summary,” Federal Highway Administration, Tech. Rep., 2010.

L. Rodegerdts, M. Blogg, E. Wemple, E. Myers, M. Kyte, M. Dixon, G. List, A. Flannery, R. Troutbeck, W. Brilon, N. Wu, B. Persaud, C. Lyon, D. Harkey, and D. Carter, “Report 572: Roundabouts in the United States,” Federal Highway Administration, Washington, D.C., Tech. Rep. 572, 2007.

L. Rodegerdts, J. Bansen, C. Tiesler, J. Knudsen, E. Myers, M. Johnson, M. Moule, B. Persaud, C. Lyon, S. Hallmark, H. Isebrands, R. B. Crown, B. Guichet, and A. O’Brien, “Report 672: Roundabouts: An Informational Guide,” Federal Highway Administration, Washington, D.C., Tech. Rep. 672, 2010.

J. A. Rodríguez-Serrano and S. Singh, “Trajectory clustering in CCTV traffic videos using probability product kernels with hidden Markov models,” *Pattern Analysis and Applications*, vol. 15, no. 4, pp. 415–426, jun 2012. [Online]. Available: <http://dx.doi.org/10.1007/s10044-012-0269-7>

E. Rosén and U. Sander, “Pedestrian fatality risk as a function of car impact speed,” *Accident Analysis & Prevention*, vol. 41, no. 3, pp. 536–542, 2009.

SAAQ, “Rapport Annuel de Gestion 2013,” Société de l’assurance automobile du Québec, Tech. Rep., 2014.

H. Sadeq, “Automated roundabout safety analysis: diagnosis and remedy of safety problems,” Ph.D. dissertation, University of British Columbia, Mar 2013. [Online]. Available: <https://open.library.ubc.ca/cIRcle/collections/24/items/1.0073622>

- L. Sakshaug, A. Laureshyn, Å. Svensson, and C. Hydén, “Cyclists in roundabouts—Different design solutions,” *Accident Analysis & Prevention*, vol. 42, no. 4, pp. 1338–1351, jul 2010. [Online]. Available: <http://dx.doi.org/10.1016/j.aap.2010.02.015>
- S. Sankararaman, P. K. Agarwal, T. Mølhave, J. Pan, and A. P. Boedihardjo, “Model-driven matching and segmentation of trajectories,” in *Proceedings of the 21st ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems - SIGSPATIAL'13*. Association for Computing Machinery (ACM), 2013. [Online]. Available: <http://dx.doi.org/10.1145/2525314.2525360>
- N. Saunier, A. El Hussein, K. Ismail, C. Morency, J. Auberlet, and T. Sayed, “Pedestrian stride frequency and length estimation in outdoor urban environments using video sensors,” *Physics of Life Reviews*, vol. 6, pp. 176–206, 2011.
- N. Saunier, “Incidence de la régulation d’un carrefour à feux sur le risque subi par les usagers Apprentissage d’indicateurs par sélection de données dans un flux,” Ph.D. dissertation, École Nationale Supérieure des Télécommunications, Paris, France, 2005.
- N. Saunier and M. G. Mohamed, “Clustering Surrogate Safety Indicators to Understand Collision Processes,” in *Transportation Research Board Annual Meeting*, 2014, 14-2380.
- N. Saunier and T. Sayed, “Clustering Vehicle Trajectories with Hidden Markov Models Application to Automated Traffic Safety Analysis,” in *The 2006 IEEE International Joint Conference on Neural Network Proceedings*. Institute of Electrical & Electronics Engineers (IEEE), 2006. [Online]. Available: <http://dx.doi.org/10.1109/IJCNN.2006.246960>
- N. Saunier and T. Sayed, “A feature-based tracking algorithm for vehicles in intersections,” in *Canadian Conference on Computer and Robot Vision*. Québec: IEEE, June 2006. [Online]. Available: <http://dx.doi.org/10.1109/CRV.2006.3>
- N. Saunier and T. Sayed, “Probabilistic framework for automated analysis of exposure to road collisions,” *Transportation Research Record: Journal of the Transportation Research Board*, no. 2083, pp. 96–104, 2008.
- N. Saunier, T. Sayed, and C. Lim, “Probabilistic Collision Prediction for Vision-Based Automated Road Safety Analysis,” in *The 10th International IEEE Conference on Intelligent Transportation Systems*. Seattle: IEEE, oct 2007, pp. 872–878. [Online]. Available: <http://dx.doi.org/10.1109/ITSC.2007.4357793>
- N. Saunier, T. Sayed, and K. Ismail, “Large-Scale Automated Analysis of Vehicle Interactions and Collisions,” *Transportation Research Record: Journal of the Transportation*

Research Board, vol. 2147, no. -1, pp. 42–50, Dec 2010. [Online]. Available: <http://dx.doi.org/10.3141/2147-06>

N. Saunier, N. Mourji, and B. Agard, “Mining Microscopic Data of Vehicle Conflicts and Collisions to Investigate Collision Factors,” *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2237, pp. 41–50, dec 2011. [Online]. Available: <http://dx.doi.org/10.3141/2237-05>

N. Saunier, P. St-Aubin, S. Burns, M. Perdomo, J.-S. Bourdeau, L. F. Miranda-Moreno, Z. Patterson, and K. Ismail, “Sécurité des carrefours giratoires: Rapport final de recherche,” Polytechnique Montréal, McGill University, Concordia University, Carleton University, Montréal, QC, Tech. Rep., 2015.

T. Schreck, J. Bernard, T. Tekusova, and J. Kohlhammer, “Visual cluster analysis of trajectory data with interactive Kohonen Maps,” in *2008 IEEE Symposium on Visual Analytics Science and Technology*. IEEE, oct 2008. [Online]. Available: <http://dx.doi.org/10.1109/VAST.2008.4677350>

N. Sekiyama, K. Minoura, and T. Watanabe, “Prediction of collisions between vehicles using attainable region,” in *Proceedings of the 5th International Conference on Ubiquitous Information Management and Communication*. ACM, 2011, p. 34.

C.-q. Shao and X.-m. Liu, “Estimation of Saturation Flow Rates at Signalized Intersections,” *Discrete Dynamics in Nature and Society*, vol. 2012, pp. 1–9, 2012. [Online]. Available: <http://dx.doi.org/10.1155/2012/720474>

C.-q. Shao, J. Rong, and X.-m. Liu, “Study on the Saturation Flow Rate and Its Influence Factors at Signalized Intersections in China,” *Procedia - Social and Behavioral Sciences*, vol. 16, pp. 504–514, 2011. [Online]. Available: <http://dx.doi.org/10.1016/j.sbspro.2011.04.471>

D. Shinar, “The traffic conflict technique: A subjective vs. objective approach,” *Journal of Safety Research*, vol. 15, no. 4, pp. 153–157, dec 1984. [Online]. Available: [http://dx.doi.org/10.1016/0022-4375\(84\)90046-X](http://dx.doi.org/10.1016/0022-4375(84)90046-X)

D. Shinar, “Speed and Crashes: A Controversial Topic and an Elusive Relationship,” in *Special Report 254: Managing Speed*, N. R. C. TRB, Ed., Washington D.C., 1998, ch. B, pp. 221–276.

J. Sorstedt, L. Svensson, F. Sandblom, and L. Hammarstrand, “A New Vehicle Motion Model for Improved Predictions and Situation Assessment,” *IEEE Trans.*

Intell. Transport. Syst., vol. 12, no. 4, pp. 1209–1219, dec 2011. [Online]. Available: <http://dx.doi.org/10.1109/TITS.2011.2160342>

B. Spicer, “A study of traffic conflicts at six intersections,” Transport and Road Research Laboratory, Crowthorne, Berkshire, UK, Tech. Rep. TRRL Report LR551, 1973.

P. St-Aubin, “Traffic Safety Analysis for Urban Highway Ramps and Lane-Change Bans Using Accident Data and Video-Based Surrogate Safety Measures,” Master’s thesis, McGill University, Montréal, 1 2012, master of Engineering.

P. St-Aubin, L. F. Miranda-Moreno, and N. Saunier, “Surrogate Safety Analysis at Protected Freeway Ramps Using Cross-sectional and Before-After Video Data,” in *Transportation Research Board (TRB) 91st Annual Meeting*, Transportation Research Board, Ed. Washington, D.C.: National Academy of Sciences, 2012.

P. St-Aubin, L. Miranda-Moreno, and N. Saunier, “An automated surrogate safety analysis at protected highway ramps using cross-sectional and before-after video data,” *Transportation Research Part C: Emerging Technologies*, vol. 36, pp. 284–295, Nov 2013. [Online]. Available: <http://dx.doi.org/10.1016/j.trc.2013.08.015>

P. St-Aubin, N. Saunier, L. F. Miranda-Moreno, and K. Ismail, “Use of Computer Vision Data for Detailed Driver Behavior Analysis and Trajectory Interpretation at Roundabouts,” *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2389, pp. 65–77, Nov 2013. [Online]. Available: <http://dx.doi.org/10.3141/2389-07>

P. St-Aubin, N. Saunier, and L. F. Miranda-Moreno, “Road User Collision Prediction Using Motion Patterns Applied to Surrogate Safety Analysis,,” in *Transportation Research Board (TRB) 93rd Annual Meeting*. Washington, D.C.: National Academy of Sciences, 2014.

P. St-Aubin, N. Saunier, and L. Miranda-Moreno, “Large-scale automated proactive road safety analysis using video data,” *Transportation Research Part C: Emerging Technologies*, vol. 58, Part B, pp. 363 – 379, 2015, big Data in Transportation and Traffic Engineering. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0968090X15001485>

P. St-Aubin, N. Saunier, and L. F. Miranda-Moreno, “Comparison of Various Time-to-Collision Prediction and Aggregation Methods for Surrogate Safety Analysis,” in *Transportation Research Board (TRB) 94th Annual Meeting*, Transportation Research Board, Ed. Washington, D.C.: National Academy of Sciences, 2015, p. 20.

P. St-Aubin, N. Saunier, and L. F. Miranda-Moreno, “Large-Scale Microscopic Traffic Behaviour and Safety Analysis of Québec Roundabout Design,” in *Transportation Research*

Board (TRB) 94th Annual Meeting, Transportation Research Board, Ed. Washington, D.C.: National Academy of Sciences, 2015, p. 18.

J. Strauss, L. F. Miranda-Moreno, and P. Morency, “Cyclist activity and injury risk analysis at signalized intersections: A Bayesian modelling approach,” *Accident Analysis & Prevention*, vol. 59, pp. 9 – 17, 2013. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0001457513001905>

Å. Svensson, “A method for analysing the traffic process in a safety perspective,” Ph.D. dissertation, Lund University, Lund, Sweden, 1998.

Å. Svensson and C. Hydén, “Estimating the severity of safety related behaviour,” *Accident Analysis & Prevention*, vol. 38, no. 2, pp. 379–385, mar 2006. [Online]. Available: <http://dx.doi.org/10.1016/j.aap.2005.10.009>

Å. Svensson, A. Laureshyn, T. Jonsson, H. Ardö, and A. Persson, “Collection of micro-level safety and efficiency indicators with automated video analysis,” in *3rd International Conference on Road Safety and Simulation*, Indianapolis, Ind., 2011.

A. Tarko, G. Davis, N. Saunier, T. Sayed, and S. Washington, “Surrogate Measures of Safety White Paper,” in *Surrogate Measures Of Safety White Paper*. ANB20(3) Subcommittee on Surrogate Measures of Safety, 2009, p. 13.

C. Tomasi and T. Kanade, “Detection and tracking of point features,” Carnegie Mellon University, Tech. Rep. CMU-CS-91-132, April 1991.

Transport Canada, *Canadian Motor Vehicle Traffic Collision Statistics 2013*. Ottawa: Transport Canada, 2015.

A. van Poortvliet, “Risks, disasters and management - A comparative study of three passenger transport systems,” Ph.D. dissertation, Technical University Delft, Netherlands, 1999.

M. Vlachos, G. Kollios, and D. Gunopulos, “Elastic Translation Invariant Matching of Trajectories,” *Machine Learning*, vol. 58, no. 2-3, pp. 301–334, 2005. [Online]. Available: <http://dx.doi.org/10.1007/s10994-005-5830-9>

N. J. Ward, J. Otto, and J. Linkenbach, “A Primer for Traffic Safety Culture,” *Institute of Transportation Engineers. ITE Journal*, vol. 84, no. 5, p. 41, 2014.

M. Williams, “Validity of the traffic conflicts technique,” *Accident Analysis & Prevention*, vol. 13, no. 2, pp. 133–145, Jun 1981. [Online]. Available: [http://dx.doi.org/10.1016/0001-4575\(81\)90025-7](http://dx.doi.org/10.1016/0001-4575(81)90025-7)

World Health Organization, *Global status report on road safety 2013*. Geneva, Switzerland: Department of violence and injury prevention and disability, 2013.

S. Zangenehpour, L. F. Miranda-Moreno, and N. Saunier, "Automated classification based on video data at intersections with heavy pedestrian and bicycle traffic: Methodology and application," *Transportation Research Part C: Emerging Technologies*, vol. 56, pp. 161–176, jul 2015. [Online]. Available: <http://dx.doi.org/10.1016/j.trc.2015.04.003>

L. Zheng, K. Ismail, and X. Meng, "Traffic conflict techniques for road safety analysis: open questions and some insights," *Canadian Journal of Civil Engineering*, vol. 41, no. 7, pp. 633–641, jul 2014. [Online]. Available: <http://dx.doi.org/10.1139/cjce-2013-0558>

L. Zheng, K. Ismail, and X. Meng, "Shifted Gamma-Generalized Pareto Distribution model to map the safety continuum and estimate crashes," *Safety Science*, vol. 64, pp. 155–162, Apr 2014. [Online]. Available: <http://dx.doi.org/10.1016/j.ssci.2013.12.003>

APPENDIX A THESIS PUBLICATIONS

The following is a list of publications and presentations that were produced as a result of this work.

A.1 Journals

- FU, Ting, ZANGENEHPOUR, Sohail, **ST-AUBIN**, Paul G., FU, Liping, MIRANDA-MORENO, Luis F. (2015) Using microscopic video data measures for driver behavior analysis during adverse winter weather: opportunities and challenges, *Journal of Modern Transportation*, vol. 53, no. 2, pp. 81-92
- **ST-AUBIN**, Paul G., SAUNIER, Nicolas, MIRANDA-MORENO, Luis F., (2015) Large-Scale Automated Proactive Road Safety Analysis Using Video Data, *Transportation Research Part C: Emerging Technologies, Special Issue: Big Data in Transportation and Traffic Engineering*, vol. 58, pp. 363-379
- **ST-AUBIN**, Paul, MIRANDA-MORENO, Luis F., SAUNIER, Nicolas (2013) An Automated Surrogate Safety Analysis at Protected Highway Ramps Using Cross-Sectional and Before-After Video Data, *Transportation Research Part C: Emerging Technologies*, vol. 36, pp. 284-295
- **ST-AUBIN**, Paul, SAUNIER, Nicolas, MIRANDA-MORENO, Luis F., ISMAIL, Karim (2013) Use of Computer Vision Data for Detailed Driver Behavior Analysis and Trajectory Interpretation at Roundabouts, *Transportation Research Record*, no. 2389, pp 65-77
- JACKSON, Stewart, MIRANDA-MORENO, Luis Fernando, **ST-AUBIN**, Paul, SAUNIER, Nicolas (2013) A Flexible, Mobile Video Camera System and Open Source Video Analysis Software for Road Safety and Behavioural Analysis, *Transportation Research Record*, no. 2365, pp 90-98

A.2 Peer-Reviewed Conference Proceedings

- MORSE, Philip, **ST-AUBIN**, Paul, MIRANDA-MORENO, Luis F., SAUNIER, Nicolas (2016) Transferability Study of Video Tracking Optimization for Traffic Data Collection and Analysis, *Proc. TRB-XCV*, Washington, D.C.; January 10-14

- **ST-AUBIN**, Paul, SAUNIER, Nicolas, MIRANDA-MORENO, Luis F. (2015) Large-Scale Microscopic Traffic Behaviour and Safety Analysis of Québec Roundabout Design, Proc. TRB-XCIV, Washington, D.C.; January 11-15
- **ST-AUBIN**, Paul, SAUNIER, Nicolas, MIRANDA-MORENO, Luis F. (2015) Comparison of Various Time-to-Collision Prediction and Aggregation Methods for Surrogate Safety Analysis, Proc. TRB-XCIV, Washington, D.C.; January 11-15
- **ST-AUBIN**, Paul, SAUNIER, Nicolas, MIRANDA-MORENO, Luis F. (2014) Big Brother is Watching You... To Predict Crashes, Big Data and Urban Informatics, Chicago, IL.; August 11-12
- **ST-AUBIN**, Paul, SAUNIER, Nicolas, MIRANDA-MORENO, Luis F. (2014) Road User Collision Prediction Using Motion Patterns Applied to Surrogate Safety Analysis, Proc. TRB-XCII, Washington, D.C.; January 12-16

A.3 Technical Reports

- SAUNIER, Nicolas, **ST-AUBIN**, Paul, BURNS, Shaun, PERDOMO, Mario, BOURDEAU, Jean-Simon, MIRANDA-MORENO, Luis F., PATTERSON, Zachary, ISMAIL, Karim (2015). Sécurité des carrefours giratoires: Rapport final de recherche. Technical report, Polytechnique Montréal, McGill University, Concordia University, Carleton University, February, 160 pages

A.4 Presentations

- **ST-AUBIN**, Paul, SAUNIER, Nicolas, MIRANDA-MORENO, Luis F. (2014) Surrogate Safety and Behavioural Analysis of Roundabout Merging Zones Using Computer Vision, 27th ICTCT Workshop, Karlsruhe, Germany; October 16-17
- GAUTHIER, Laurent, **ST-AUBIN**, Paul, SAUNIER, Nicolas (2014) Détection des mouvements potentiellement dangereux dans un rond-point, Colloque 2014 du RRSR, Québec; April 24
- **ST-AUBIN**, Paul, SAUNIER, Nicolas, MIRANDA-MORENO, Luis F. (2014) Analyse du comportement des conducteurs Québécois et de leurs interactions aux carrefours giratoires, 49e congrès annuel de l'AQTR, Québec; April 1

- **ST-AUBIN**, Paul (2013) Application de la vision par ordinateur pour le calcul d'indicateurs substitutifs de sécurité routière pour les aménagements routiers, Colloque 2013 du RRSR, Québec; May 3

APPENDIX B SOFTWARE

B.1 Data Storage

Metadata is stored as a SQLite database in the root folder (`config.dir`). This file organizes and stores the file structure of all video sequences following the metadata specification of section 3.4.1. It is used to point to individual video sequences or to run specific analyses. Individual video sequences are stored in a schema-compatible structure as follows:

```
config.dir / site.name / camera.name / sequence.name
```

As an example, a database is prepared on drive `config.dir = I:`. A 20-minute video sequence taken at 9:40 a.m. on the 12th of June 2012 from the second camera at the intersection of *Chemin du Golf* and *Boulevard de l'Île des Soeurs* would be stored using the following path:

```
I:\MTG-desSoeurs-duGolf\V2\20120712\0940.avi
```

where `sequence.name` can be set to `20120712\0940.avi`.

B.2 Software Stack

The following software was used. Versions are as of writing.

- **Python** 2.7 64-bit (Python) is a general purpose, high-level programming language for rapid prototyping and scripting tasks.¹
- **SciPy stack** (Python) is a collection of scientific libraries expanding Python, including **SciPy** 0.16, **Numpy** 1.1, and **Matplotlib** 1.5.²
- **SQLAlchemy** 1.0 (Python) is a Python framework for working with SQLite databases.³
- **munkres** 1.0.7 (Python) is a Python library implementing the Kuhn-Munkres algorithm for solving the Assignment Problem. This library is required for MOTA analysis, but is otherwise optional.⁴
- **Python Imaging Library (PIL)** 1.1.7 (Python) is a python image library required for some functions of Traffic Intelligence.⁵

¹<https://www.python.org/>

²<https://www.scipy.org/>

³<http://www.sqlalchemy.org/>

⁴<http://software.clapper.org/munkres/>

⁵<http://www.pythonware.com/products/pil/>

- **MEncoder** 6.1 is an optional, free video transcoding tool.⁶
- **Trajectory Management and Analysis** rev 61 (C++) is a framework for storing trajectory data.⁷
- **OpenCV** 3.1 (C++, Python) is an open-source computer vision framework.⁸
- **Traffic Intelligence** rev 758 (C++, Python) applies OpenCV to traffic video data, stores trajectory data, and includes basic analysis tools. It depends on Trajectory Management and Analysis.⁹
- **tvaLib** 2.2.1.6 rev 356 (Python) is a collection of tools to extend and enhance Traffic Intelligence analysis functionality and to automate a large variety of analysis tasks.¹⁰
- **Urban Tracker Annotation Tool** commit 9ba9415 (C++) is an optional tool used to annotate video manually.¹¹

B.2.1 Software Functionality

tvaLib includes the following functionality, in typical order of usage:

- Folder-based video file concatenation into hour-long sequences.
- GUI-assisted scene drawing tools to quickly build video metadata, including homographies (GUI ported from Traffic Intelligence), masks, analysis zones, trajectory-space transformation, alignments, HLI elements, and more.
- A distortion matrix calculator and undistortion tool, as per section 3.2.2.
- MOTA (Bernardin and Stiefelhagen, 2008) calculator using `computeClearMOT()` from the `moving` module of Traffic Intelligence, and tracking configuration optimizer using MOTA and a rudimentary genetic algorithm.
- A Traffic Intelligence wrapper to automate, using metadata, feature tracking jobs in parallel (from an asynchronous pool).
- Object data loading and caching. See section 3.4.4 for more information.

⁶<https://www.mplayerhq.hu/DOCS/HTML/en/mencoder.html>

⁷<https://bitbucket.org/trajectories/trajectorymanagementandanalysis>

⁸<http://opencv.org/>

⁹<https://bitbucket.org/Nicolas/trafficintelligence>

¹⁰<https://bitbucket.org/pstaub/tvaLib>

¹¹<http://www.jpjodoin.com/urbantracker/>

- Tracking data clean-up and object filtering routines (see section 3.4.6 for specifics). These filtering routines can be calibrated using the MOTA and genetic algorithm optimization functionality listed above.
- A machine learning pass to prepare any analysis methodologies that require training or calibration using existing data, such as may be necessary with motion pattern path prediction (see section 4.3.2).
- Interaction (conflict) analysis and caching functions.
- A basic statistical traffic report and basic analysis figures (speed distributions, conflicts, geometry, etc.).
- A report compiler (PDF LaTeX format) with support for HLI data.

Furthermore, a number of scripts have been prepared to batch process some of the above tasks—with an option to parallelize tasks using an asynchronous pool—and to conduct overall summaries and statistical analyses on a site basis. This includes, among others:

- A batch processing script which runs the above functions in parallel over multiple sequences simultaneously.
- An analysis script which crawls serialized data to generate overall reports.
- An HLI analysis script with similar functionality, but operating specifically on HLI modules.
- A PDF report compiler.

B.3 Code

tvaLib heavily uses Traffic Intelligence libraries, functionality, and data storage specifications. The command `-h` when running `main.py` lists all available commands.

B.3.1 Style

Mathematical variables use underscores, and other variables (including class instances), and function/method names use camelCase (with a lower-case first letter), while class definitions use CamelCase (with an upper-case first letter). Lists, arrays, and dicts have an “s” appended to the variable name, while indices, usually a single letter (i.e. “sIx”), have “ix” appended to them. “get”, “set”, and “gen” should prefix most common class methods.

B.3.2 tvaLib File Structure

The `lib` contains the bulk of the library stored in the form of various functions and classes. In the root folder, `main.py` (or `main.exe` or `main` if using a compiled build for Windows or Unix respectively) is used to demonstrate implementation of these tools and launch annotation, filtering, caching, core computations, sequence analysis, and plotting of most tracking data. The `scripts` folder contains processing scripts used to automate `main.py` batch tasks and to run other data compilation tasks. The `include` folder contains runtime functions which are used exclusively by `main.py`. The `h1i` folder extends `tvaLib` functionality for application-specific code (see section B.3.4).

B.3.3 Accessing `movingObject()` Data

The following code B.1 demonstrates access to trajectory data from an instance of class `MovingObject()` in `Traffic Intelligence` and `tvaLib`. Lines 1 and 2 are used to obtain a sequence (list) of `x` and `y` coordinates respectively for the moving object, in standard units of distance. Lines 3 and 4 are used to obtain a sequence (list) of `x` and `y` components of the sequence of velocity vectors, in standard units of distance per frame.

Listing B.1 `MovingObject()` common methods

```
obj.getXCoordinates()           # [ 35.67,  34.65, ...,
p_x_n]
obj.getYCoordinates()          # [101.34, 103.04, ...,
p_y_n]
obj.velocities.getXCoordinates() # [ -1.02, ...,
v_x_n-1, v_x_n-1]
obj.velocities.getYCoordinates() # [  1.70, ...,
v_y_n-1, v_y_n-1]
obj.curvilinearPositions.getXCoordinates() # [ 25.01, 26.99, ...,
S_n]
obj.curvilinearPositions.getYCoordinates() # [  0.60,  0.67, ...,
Y_n]
obj.curvilinearPositions.getLanes() # [  1,      1,      ...,
lane_n]
obj.getFirstInstant()           # 3475 (in frames)
obj.getLastInstant()           # 3475+n (in frames)
```

B.3.4 High-Level Interpretation Modules

To add High-Level Interpretation (HLI) functionality to tvaLib, a HLI-compatible module (Python script) should be added to the `hli` folder and tvaLib will run the `main()` function in the module automatically at the end of the ordinary analysis when using the command `--hli HLI_NAME` or alternatively by running all HLI modules using the command `--hli all`. The benefit of HLI is that these modules will inherit all commands, configurations, objects, interactions, and metadata from the main program, preloaded and pre-treated. These modules can be developed and installed independantly from tvaLib or other HLI modules. HLI modules are a type of software plugin.

APPENDIX C STATISTICS

Table C.1 lists the commands used in Stata 12.0 to perform the regression analysis and associated tests.

Table C.1 Stata Regression Commands and PostEstimation Statistics

Command	Description
<code>regress</code>	Perform ordinary least squares linear regression
<code>xtreg</code>	Perform linear regression on panel data (fixed, between, and random effects)
<code>xttest0</code>	Breusch and Pagan Lagrangian multiplier test of variance for random effects
<code>test</code>	Wald test (automatically performed with most regression models)
<code>hausman</code>	Hausman's specification test
<code>swilk</code>	Shapiro-Wilk W-test for normality of residuals
<code>estat hettest</code>	Breusch and Pagan Lagrangian test for homoscedasticity of residuals
<code>linktest</code>	Model specification link test for single-equation models
<code>ovtest</code>	Ramsey RESET test of regression model specification

APPENDIX D COMPLEMENTARY DATA

D.1 Site Selection Score

Site selection is performed by sorting sites on the basis of a ten-point score, averaging a number of measures (each on a ten-point score) evaluating data collection suitability. This suitability reflects a number of needs including minimum number of traffic observations to be included in the study, data collection costs, and equipment installation feasibility. These are summarized in Table D.1. Furthermore, to ensure good statistical sampling, sites with regional, administrative, and land use characteristics identical to those of another, but having an inferior score, are automatically rejected, even if these have better scores than other sites.

Table D.1 Site Selection Scale Parameters

Measure	Description
Estimated traffic	Traffic data is not available for every single intersection the entire province of Québec. Subjective estimation on the basis of local land use and network type is performed.
Clearance (installation)	A rating performed via visual inspection of recent site photographs which determines suitability of installing VCUs and dealing with installation issues such as power lines, availability of lamp posts, etc.
Clearance (visibility)	A rating performed by visual inspection of recent site photographs which determines suitability of installing VCUs and dealing with issues of visibility and visual obstruction (e.g. trees).
Installation equipment	A rating performed via visual inspection of recent site photographs which determines suitability of installing VCUs and dealing with issues of equipment compatibility.
Travel expenditure	A rating based on the <i>traveldistance</i> (in km) necessary to perform data collection, following: $score = \frac{10}{e^{\frac{0.003traveldistance}{\min(3,n)}}}$ where n is the number of sites in the vicinity of the destination to be instrumented simultaneously.

D.2 Time-to-Collision Regression Fit

Figure D.1 plots the fit of all 41 probability distributions fit in section 5.3.4.




Figure D.1 Curve fitting of all TTC_{15} values for the 41 best probability density function models across all sites.