

Driver Behavior in Traffic

Publication No. FHWA-HRT-12-036

February 2012



U.S. Department of Transportation
Federal Highway Administration

ACKNOWLEDGEMENTS

The authors would like to express thanks to Dr. C. Y. David Yang, FHWA's Agreement Officer's Technical Representative, for his continued support and guidance during this project. They would also like to thank individuals at Virginia Tech and the Virginia Tech Transportation Institute who contributed to the study in various ways: Greg Fitch, Shane McLaughlin, Brian Daily, Clark Gaylord, Rich Hanowski, Myra Blanco, Kelly Stanley, and Rebecca Olson.

Notice

This document is disseminated under the sponsorship of the U.S. Department of Transportation in the interest of information exchange. The U.S. Government assumes no liability for the use of the information contained in this document. This report does not constitute a standard, specification, or regulation.

The U.S. Government does not endorse products or manufacturers. Trademarks or manufacturers' names appear in this report only because they are considered essential to the objective of the document.

Quality Assurance Statement

The Federal Highway Administration (FHWA) provides high-quality information to serve Government, industry, and the public in a manner that promotes public understanding. Standards and policies are used to ensure and maximize the quality, objectivity, utility, and integrity of its information. FHWA periodically reviews quality issues and adjusts its programs and processes to ensure continuous quality improvement.

Technical Report Documentation Page

1. Report No. FHWA-HRT -12 -036	2. Government Accession No.	3. Recipient's Catalog No.	
4. Title and Subtitle Driver Behavior in Traffic		5. Report Date February 2012	
		6. Performing Organization Code:	
7. Author(s) Montasir Abbas, Alejandra Medina, Linsen Chong, Bryan Higgs, Catherine McGhee, Mike Fontain, Noah Goodall, Lukas Kautzsch, Axel Leonhardt, and Kiel Ova		8. Performing Organization Report No.	
9. Performing Organization Name and Address Virginia Polytechnic Institute and State University Office of Sponsored Programs 1880 Pratt Drive, Suite #2006 Blacksburg, VA 24060		10. Work Unit No.	
		11. Contract or Grant No. DTFH61-09-H-00007	
12. Sponsoring Agency Federal Highway Administration Turner-Fairbank Highway Research Center Exploratory Advanced Research Program 6300 Georgetown Pike McLean, VA 22101-2296		13. Type of Report and Period Covered "[Final Report]"	
		14. Sponsoring Agency Code	
15. Supplementary Notes AOTR – Dr. C. Y. David Yang, Office of Operations R&D			
16. Abstract Existing traffic analysis and management tools do not model the ability of drivers to recognize their environment and respond to it with behaviors that vary according to the encountered driving situation. The small body of literature on characterizing drivers' behavior is typically limited to specific locations (i.e., by collecting data on specific intersections or freeway sections) and is very narrow in scope. This report documented the research performed to model driver behavior in traffic under naturalistic driving data. Of special interest to this research was the modeling of both normal and safety-critical driving. The research resulted in the development of a hybrid car-following model. In addition, neuro-fuzzy reinforcement learning, an agent-based artificial intelligence machine-learning technique, was used to model driving behavior. The naturalistic driving database was used to train and validate driver agents. The proposed methodology simulated events from different drivers and proved behavior heterogeneities. Robust agent activation techniques were also developed using discriminant analysis. The developed agents were implemented in the VISSIM simulation platform and were evaluated by comparing the behavior of vehicles with and without agent activation. The results showed very close resemblance of the behavior of agents to driver data. Prototype agents (spreadsheets and codes) were developed. Future research recommendations include training agents using more data to cover a wider region in the Wiedemann regime space, and sensitivity analysis of agent training parameters. The goal of this effort is to provide the industry with methods for developing more accurate and more sensitive traffic simulation models. This could also enable future research to develop new generations of traffic simulation models that accurately model driver behavior during incidents and other complex traffic situations.			
17. Key Words Naturalistic Driving, Agent-Based Modeling and Simulation, Reinforcement Learning, Artificial Intelligence, Neuro-Fuzzy Actor-Critic		18. Distribution Statement No restrictions. This document is available through the National Technical Information Service; Springfield, VA 22161.	
19. Security Classif. (of this report) Unclassified	20. Security Classif. (of this page) Unclassified	21. No. of Pages 160	22. Price

SI* (MODERN METRIC) CONVERSION FACTORS

APPROXIMATE CONVERSIONS TO SI UNITS

Symbol	When You Know	Multiply By	To Find	Symbol
LENGTH				
in	inches	25.4	millimeters	mm
ft	feet	0.305	meters	m
yd	yards	0.914	meters	m
mi	miles	1.61	kilometers	km
AREA				
in ²	square inches	645.2	square millimeters	mm ²
ft ²	square feet	0.093	square meters	m ²
yd ²	square yard	0.836	square meters	m ²
ac	acres	0.405	hectares	ha
mi ²	square miles	2.59	square kilometers	km ²
VOLUME				
fl oz	fluid ounces	29.57	milliliters	mL
gal	gallons	3.785	liters	L
ft ³	cubic feet	0.028	cubic meters	m ³
yd ³	cubic yards	0.765	cubic meters	m ³
NOTE: volumes greater than 1000 L shall be shown in m ³				
MASS				
oz	ounces	28.35	grams	g
lb	pounds	0.454	kilograms	kg
T	short tons (2000 lb)	0.907	megagrams (or "metric ton")	Mg (or "t")
TEMPERATURE (exact degrees)				
°F	Fahrenheit	5 (F-32)/9 or (F-32)/1.8	Celsius	°C
ILLUMINATION				
fc	foot-candles	10.76	lux	lx
fl	foot-Lamberts	3.426	candela/m ²	cd/m ²
FORCE and PRESSURE or STRESS				
lbf	poundforce	4.45	newtons	N
lbf/in ²	poundforce per square inch	6.89	kilopascals	kPa

APPROXIMATE CONVERSIONS FROM SI UNITS

Symbol	When You Know	Multiply By	To Find	Symbol
LENGTH				
mm	millimeters	0.039	inches	in
m	meters	3.28	feet	ft
m	meters	1.09	yards	yd
km	kilometers	0.621	miles	mi
AREA				
mm ²	square millimeters	0.0016	square inches	in ²
m ²	square meters	10.764	square feet	ft ²
m ²	square meters	1.195	square yards	yd ²
ha	hectares	2.47	acres	ac
km ²	square kilometers	0.386	square miles	mi ²
VOLUME				
mL	milliliters	0.034	fluid ounces	fl oz
L	liters	0.264	gallons	gal
m ³	cubic meters	35.314	cubic feet	ft ³
m ³	cubic meters	1.307	cubic yards	yd ³
MASS				
g	grams	0.035	ounces	oz
kg	kilograms	2.202	pounds	lb
Mg (or "t")	megagrams (or "metric ton")	1.103	short tons (2000 lb)	T
TEMPERATURE (exact degrees)				
°C	Celsius	1.8C+32	Fahrenheit	°F
ILLUMINATION				
lx	lux	0.0929	foot-candles	fc
cd/m ²	candela/m ²	0.2919	foot-Lamberts	fl
FORCE and PRESSURE or STRESS				
N	newtons	0.225	poundforce	lbf
kPa	kilopascals	0.145	poundforce per square inch	lbf/in ²

*SI is the symbol for the International System of Units. Appropriate rounding should be made to comply with Section 4 of ASTM E380.
(Revised March 2003)

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	2
EXECUTIVE SUMMARY	XIII
OVERVIEW	XIII
RESEARCH OBJECTIVES	XIII
NATURALISTIC DRIVING DATA.....	XIII
EVALUATION AND EXTENSION OF EXISTING MODELS	XV
AGENT DEVELOPMENT	XVI
SIMULATION AND ANALYSIS	XX
<i>Summary of Operations Analysis</i>	xx
<i>Summary of Safety Surrogate Analysis Model</i>	xxi
CONCLUSIONS AND RECOMMENDATIONS	XXII
CHAPTER 1. INTRODUCTION.....	1
BACKGROUND	1
OVERALL SCOPE OF THE EFFORT.....	1
REPORT ORGANIZATION	2
CHAPTER 2. LITERATURE REVIEW	3
INTRODUCTION	3
CAR-FOLLOWING MODELS	3
<i>Pipes Model</i>	4
<i>The GM models</i>	5
<i>The Wiedemann Model</i>	7
<i>The Fritzsche Model</i>	9
<i>The CA Model</i>	11
<i>The SK Model</i>	14
<i>The IDM Model</i>	15
<i>The Gipps Model</i>	16
<i>The VDIFF Model</i>	17
CAR-FOLLOWING MODELS AND SIMULATION SOFTWARE	17
MODELING DRIVER BEHAVIOR.....	18
MODEL CALIBRATION WITH LONGITUDINAL TRAJECTORY DATA	19
DRAWBACKS OF PARAMETER CALIBRATION METHODS	20
ALTERNATE MODELING OF THE DRIVING PROCESS	20
DRIVING BEHAVIOR HETEROGENEITY	20
INTRODUCING INCIDENT PARAMETERS WITH TRAJECTORY DATA.....	21
NATURALISTIC DATA	22
<i>Naturalistic Data: 100-Car Database</i>	22
<i>Naturalistic Data: 34-Truck Database</i>	23
<i>Naturalistic Data: 8-Truck Database</i>	23
AGENT-BASED MODELING IN TRANSPORTATION	25
REINFORCEMENT LEARNING METHOD.....	26
REINFORCEMENT LEARNING USED IN TRAFFIC ADAPTIVE CONTROL SYSTEMS	27
REINFORCEMENT LEARNING IN TRANSPORTATION NETWORK PROBLEMS	28
REINFORCEMENT LEARNING APPLIED IN TRAFFIC KINEMATIC PROBLEM.....	29
CHAPTER 3. NATURALISTIC DATA PROCESSING	33
NATURALISTIC TRUCK DATABASE AND TRIGGERS.....	37
NATURALISTIC CAR DATABASE	39
SAFETY-CRITICAL EVENTS	40
CAR-FOLLOWING FILTERING CRITERIA	43
CHAPTER CONCLUSION	45

CHAPTER 4. CALIBRATION AND EXTENSION OF EXISTING CAR-FOLLOWING MODELS—THE WIEDEMANN MODEL	46
MODEL CALIBRATION.....	46
<i>Wiedemann Car-Following Model</i>	46
<i>Reconstruction of the Wiedemann Model</i>	47
<i>Calibration and Adaptation of Existing Model</i>	47
<i>New Thresholds</i>	52
RESULTS	53
<i>Evaluation of the Existing Wiedemann Model</i>	53
<i>New Thresholds</i>	55
<i>Reconstructed Wiedemann Model</i>	56
<i>Conclusions</i>	59
HYBRID MODEL DEVELOPMENT	59
<i>GHR Model</i>	59
<i>Wiedemann Model Equations</i>	59
<i>Calibration Framework</i>	61
<i>Results</i>	61
<i>Conclusions</i>	64
MODEL COMPARISON	64
<i>Conclusions</i>	70
CHAPTER CONCLUSION	70
CHAPTER 5. AGENT DEVELOPMENT AND TRAINING WITH REINFORCEMENT LEARNING	72
BACKGROUND INFORMATION	72
TRAFFIC STATES AND ACTIONS DURING SAFETY-CRITICAL EVENTS	72
AGENT-BASED MODELING OF DRIVER BEHAVIOR	72
POTENTIAL SHORTCOMINGS OF USING EXISTING REINFORCEMENT LEARNING ALGORITHMS FOR AGENT TRAINING	73
THE PROPOSED REINFORCEMENT LEARNING METHODOLOGY	73
AGENT TRAINING METHODOLOGY: NFACTRL.....	74
<i>NFACTRL Structure</i>	74
<i>Weights Update: Reinforcement Learning Algorithm</i>	80
NATURALISTIC DATA EXTRACTION AND AGENT TRAINING PROCEDURE	82
<i>Database Description</i>	82
<i>Safety-Critical Events Extraction</i>	82
<i>State and Action Variables Selection</i>	83
<i>Training Process</i>	83
AGENT TRAINING RESULTS.....	84
CROSS VALIDATION RESULTS	87
THE MEGA-AGENT CONCEPT	91
CHAPTER CONCLUSION.....	95
CHAPTER 6. SAFETY-RELATED AGENT ACTIVATION	97
DISCRIMINANT ANALYSIS.....	97
<i>Car Drivers</i>	97
<i>Truck Drivers</i>	110
CHAPTER CONCLUSION	119
CHAPTER 7. AGENT IMPLEMENTATION IN VISSIM.....	121
INTRODUCTION AND PURPOSE	121
OVERALL WORKFLOW.....	121
VISSIM SETTINGS.....	123
<i>Vehicle Types</i>	123
<i>2D/3D Model Distribution</i>	123
DRIVING BEHAVIOR	124

ACCELERATION AND DECELERATION FUNCTIONS	124
PARAMETER FILE *.APC	125
<i>GUI</i>	125
<i>Parameter Contents and Format</i>	126
SCENARIO DESCRIPTION "EVENT"	127
<i>Description</i>	127
<i>How to Run</i>	128
CHAPTER CONCLUSION	135
CHAPTER 8. SIMULATION AND ANALYSIS.....	136
INTRODUCTION	136
<i>Simulation Environment</i>	136
<i>Evaluation</i>	136
OPERATIONS ASSESSMENT	136
<i>Agent Behavior in VISSIM</i>	136
<i>Quantitative Assessment of Vehicle Behavior</i>	141
<i>Summary of Operations Analysis</i>	144
SAFETY ASSESSMENT	145
<i>General Observations on Agent Behavior and Impact on Safety Surrogates</i>	146
EVALUATION OF AGENTS ON MORE COMPLEX NETWORKS.....	152
CHAPTER CONCLUSION	153
CHAPTER 9. CONCLUSIONS AND RECOMMENDATIONS	155
CONCLUSIONS	155
RESEARCH CONTRIBUTION	155
RECOMMENDATION FOR FUTURE RESEARCH.....	156
REFERENCES	157

LIST OF FIGURES

Figure 1. Screen shot. Naturalistic driving data collection.	xiv
Figure 2. Graph. Comparison of car-following models.	xvi
Figure 3. Illustration. Agent learning process.	xvii
Figure 4. Graph. Acceleration of mega-agent.	xviii
Figure 5. Chart. Yaw angle of mega-agent.	xix
Figure 6. Illustration. System representation of driver’s response to encountered traffic situation.	2
Figure 7. Graph. Wiedemann car-following model’s thresholds (AX, ABX, SDX, CLDV, SDV, OPDV).⁽¹⁷⁾	7
Figure 8. Graph. Different thresholds and regimes in the Fritzsche car-following model.⁽³⁾	10
Figure 9. Illustration. “Bottom-up” modeling of system.⁽⁵⁹⁾	26
Figure 10. Diagram. Advantages of the naturalistic data collection methodology data collection and reduction.⁽⁴⁵⁾	33
Figure 11. Illustration. Naturalistic data sensors.⁽⁴⁸⁾	34
Figure 12. Screen Shot. Screen shot of DART software interface.	35
Figure 13. Illustration. Truck driver demographic data.	42
Figure 14. Illustration. Car driver demographic data.	43
Figure 15. Diagram. Wiedemann 74 car-following logic.⁽⁷³⁾	46
Figure 16. Graph. Example hook car-following period.	53
Figure 17. Graph. Hooked/not-hooked division line.	55
Figure 18. Graph. Pass decision curve.	56
Figure 19. Graph. Thresholds for Driver J.	57
Figure 20. Graph. Thresholds for Driver K.	57
Figure 21. Graph. Thresholds for Driver F.	58
Figure 22. Graph. Thresholds for Driver L.	58
Figure 23. Graph. Driver F Wiedemann car-following period.	61
Figure 24. Graph. Driver F Wiedemann with GHR car-following period.	62
Figure 25. Graph. Comparison of Wiedemann and Wiedemann with GHR.	62
Figure 26. Graph. Comparison between models in range versus range rate.	63
Figure 27. Graph. Plot of models as compared to the data for Driver L.	70
Figure 28. Diagram. NFACRL structure.	75

Figure 29. Graph. Acceleration of Agent A, event A1.....	84
Figure 30. Graph. Yaw angle of Agent A, event A1.....	85
Figure 31. Graph. Acceleration of Agent A, event A2.....	86
Figure 32. Graph. Yaw angle of Agent A, event A2.....	86
Figure 33. Graph. Acceleration of Agent B, event B1.....	88
Figure 34. Graph. Yaw angle of Agent B, event B1.....	88
Figure 35. Graph. Acceleration of Agent B, event A1.....	89
Figure 36. Graph. Yaw angle of Agent B, event A1.....	89
Figure 37. Graph. Acceleration of Agent A, event B1.....	90
Figure 38. Graph. Acceleration of Agent A, event B1.....	90
Figure 39. Graph. Acceleration of mega-agent, event A1.....	92
Figure 40. Graph. Yaw angle of mega-agent, event A1.....	92
Figure 41. Graph. Acceleration of mega-agent, event B1.....	93
Figure 42. Graph. Yaw angle of mega-agent, event B1.....	93
Figure 43. Graph. Percentage of misclassified datapoints for Driver l.....	98
Figure 44. Graph. Discriminant scores for Driver l.....	99
Figure 45. Graph. Longitudinal acceleration term for Driver 1.....	99
Figure 46. Graph. Range term for Driver 1.....	100
Figure 47. Graph. Range rate term for Driver 1.....	100
Figure 48. Graph. Discriminant scores of a sample car-following period.....	101
Figure 49. Graph. Discriminant scores of a sample near-crash event.....	102
Figure 50. Graph. Trajectory of a sample near-crash event.....	102
Figure 51. Graph. Percent misclassification for Driver m.....	104
Figure 52. Graph. Discriminant scores for Driver m.....	104
Figure 53. Graph. Longitudinal acceleration term for Driver m.....	105
Figure 54. Graph. Range term for Driver m.....	105
Figure 55. Graph. Range rate term for Driver m.....	106
Figure 56. Graph. Discriminant scores of a sample car-following period.....	107
Figure 57. Graph. Discriminant scores of a sample near-crash event.....	107
Figure 58. Graph. Trajectory of a sample near-crash event with the 7-s crash prediction point highlighted.....	108
Figure 59. Graph. Discriminant score for using Driver 1’s coefficients for Driver m.....	108
Figure 60. Graph. Discriminant scores for using Driver m’s coefficients for Driver l.....	109

Figure 61. Graph. Percentage of misclassified datapoints for Driver E.....	111
Figure 62. Graph. Discriminant scores for Driver E.....	111
Figure 63. Graph. Longitudinal acceleration term for Driver E.....	112
Figure 64. Graph. Range term for Driver E.....	112
Figure 65. Graph. Range rate term for Driver E.....	113
Figure 66. Graph. Percentage of misclassification for Driver G.	114
Figure 67. Graph. Discriminant scores for Driver G.	115
Figure 68. Graph. Longitudinal acceleration term for Driver G.	116
Figure 69. Graph. Range term for Driver G.	116
Figure 70. Graph. Range rate term for Driver G.	117
Figure 71. Graph. Discriminant score for using the joint coefficients for Driver E.....	118
Figure 72. Graph. Discriminant scores for using the joint coefficients for Driver G.....	118
Figure 73. Chart. Overall workflow.....	122
Figure 74. Screen shot. Specification of vehicle types to use the agent model.....	123
Figure 75. Screen shot. Model distribution for agent vehicle type.....	123
Figure 76. Screen shot. Model distribution dialog.....	124
Figure 77. Screen shot. Lateral driving behavior.....	124
Figure 78. Screen shot. Agent parameter GUI—action discrete sets.....	126
Figure 79. Screen shot. Agent parameter GUI—acceleration weights.....	126
Figure 80. Screen shot. “Run VISSIM example” command.....	129
Figure 81. Screen shot. “Run Excel Agent based on VISSIM data” command.....	130
Figure 82. Graph. Comparison of speeds and trajectories of leading vehicle and agent over time.	131
Figure 83. Graph. Comparison of acceleration over time (acceleration calculated in Excel as reference, acceleration calculated in VISSIM and performed in VISSIM to be compared).	132
Figure 84. Graph. Comparison of lane angle over time (lane angle calculated in Excel as reference, lane angle calculated in VISSIM and performed in VISSIM to be compared).	133
Figure 85. Graph. Acceleration scatter plot with linear regression analysis, Excel agent versus VISSIM agent.....	134
Figure 86. Graph. Lane angle scatter plot with linear regression analysis, Excel agent versus VISSIM agent.....	135
Figure 87. Graphs. Time-space graphs of each simulation run of Agent g, both with the agent on and the agent off.	139
Figure 88. Chart. Wiedemann diagram of a single simulation run of Agent g.....	140

Figure 89. Chart. Wiedemann diagram of a single simulation run of Agent I. 141
Figure 90. Screenshot. Example of interchange on test network. 153

LIST OF TABLES

Table 1. R-squared values of the truck mega-agent.	xix
Table 2. R-squared values of the car mega-agent.	xix
Table 3. Major similarities and differences between car-following models.	17
Table 4. Comparison of parameter values in GM-based model for acceleration and deceleration models.⁽²⁶⁾	18
Table 5. Data Analysis and Reduction Tool (DART) components.⁽⁴⁵⁾	36
Table 6. Trigger values used in the DDWS FOT.⁽⁴⁷⁾	36
Table 7. Trigger values used in the NTDS.⁽⁴⁸⁾	39
Table 8. Trigger values used in the 100-Car Study.⁽⁴⁵⁾	40
Table 9. Enumeration of crash and near-crash data.	41
Table 10. Random sample results for number of hook following periods.	52
Table 11. Calibration parameters by driver.....	54
Table 12. Root mean square error for optimization function.....	54
Table 13. Hook following threshold equations by driver.	55
Table 14. Pass threshold equations by driver.....	56
Table 15. GHR calibrated parameters by driver and regime.....	63
Table 16. Root mean square error by driver and model.....	64
Table 17. Root mean squared error for each model by driver.	65
Table 18. Sum of squared error for the models by driver with a root mean squared error of the total.	65
Table 19. GHR parameter values by driver.	66
Table 20. Gipps parameter values by driver.	66
Table 21. IDM parameter values by driver.	66
Table 22. VDIFF parameter values by driver.	67
Table 23. Wiedemann parameter values by driver.....	68
Table 24. Fritzsche parameter value by driver.	69
Table 25. R squared values for cross validation.....	91
Table 26. R squared values of the test mega-agent.	94
Table 27. R squared values of the truck mega-agent.....	94
Table 28. R squared valued of the car mega-agent.....	95
Table 29. Discriminant coefficients for Driver l.....	98
Table 30. Discriminant coefficients of Driver m.	103

Table 31. Driver-specific discriminant coefficients and thresholds.	109
Table 32: Discriminant coefficients for Driver E.....	110
Table 33. Discriminant coefficients of Driver G.	114
Table 34. Discriminant coefficients from joint analysis.	117
Table 35. 34-Truck driver-specific coefficients and thresholds.....	119
Table 36. 8-Truck driver-specific coefficients and thresholds.....	119
Table 37. Parameter file description.	127
Table 38. Mobility performance metrics for car agents, average values over 20 simulation runs.....	142
Table 39. Mobility performance metrics for truck agents, average values over 20 simulation runs.....	143
Table 40. Safety surrogate measure comparison using 1.5-s TTC – car agents.....	148
Table 41. Safety surrogate measure comparison using 2.5-s TTC – car agents.....	149
Table 42. Safety surrogate measure comparison using 1.5-s TTC – truck agents.....	150
Table 43. Safety surrogate measure comparison using 2.5-s TTC – truck agents.....	151

EXECUTIVE SUMMARY

OVERVIEW

Existing traffic analysis and management tools do not model drivers' ability to recognize their environment and respond to it differently depending on the situation. The literature on characterizing driver behavior is very limited. What research exists is typically limited to specific locations (i.e., by collecting data on specific intersections or freeway sections) and is very limited in scope. The majority of traffic modeling and parameter calibration research has assumed somewhat similar driving conditions and behavioral sets for the entire driver population. This approach does not capture or predict the effects of situational factors on individual drivers' actions.

RESEARCH OBJECTIVES

This report documents the research findings of the Federal Highway Administration Exploratory Advanced Research Program's "Driver Behavior in Traffic" study. The goal of this research was to characterize driver behavior under naturalistic driving experiences with respect to critical parameters related to freeway driving. More specifically, the research aimed to determine critical parameters related to freeway driving during critical events, including target speed selection, accelerations, lane changing behavior, car-following distance, response times, and emergency stopping behavior for representative conditions. Of special interest to this research was the integrated modeling of both normal and safety-critical driving beyond existing models' capabilities. Naturalistic driving data was analyzed and used as the basis for training 20 simulation agents, each of which encapsulated an individual drivers' decisions in response to varying traffic situations. The developed agents were designed and trained to learn individual drivers' actions for any given traffic state. State and action combinations for each individual driver were retrieved from the Virginia Tech Transportation Institute's (VTTI) database of naturalistic driving data. These characteristic driving rules of the agents were coded in a microscopic simulation environment (VISSIM) to test and study their effects on agents' behavior in different conditions and scenarios. This research is expected to provide the simulation and modeling industry with methods for developing more accurate and more sensitive traffic models. It could also enable future research to develop new traffic simulation models that accurately model driver behavior during incidents and other complex traffic situations.

NATURALISTIC DRIVING DATA

The research conducted in this project uses naturalistic data as a source for training agents to mimic the behavior of selected drivers during normal and safety-critical events. Naturalistic data collection is the collection of driver behavior and performance data in a real-world environment. Naturalistic data analysis provides an opportunity to examine what happens in the final seconds before crash, near-crash, and safety-critical events for which an analyst would otherwise depend on eye witness accounts and police reports. Data regarding vehicle speed, acceleration, range, range rate, headway, time to collision (TTC), brake pedal input, and qualitative data, such as pre-incident maneuvers, can be used to describe driver behavior. Other data such as roadway type, number of lanes, traffic density, time of day, and weather can be used to describe the driving

environment. Tying these data together allows researchers to understand the exact conditions that exist during crash events. Naturalistic driving data collection is a powerful approach used by researchers to understand crash and near-crash causation. As opposed to traditional epidemiological and experimental/empirical approaches, this *in situ* process uses drivers who operate vehicles that have been equipped with specialized sensors as well as processing and recording equipment.

Naturalistic data collection methods require a sophisticated network of sensors, processors, and recording systems. The Data Acquisition System (DAS) provides a diverse collection of both on-road driving and driver (participant, non-driving) data, including measures such as driver input and performance (e.g., lane position, headway, etc.), four camera video views, and driver activity data (Figure 1). This information may be supplemented by subjective data, such as questionnaire data. A tremendous amount of data is acquired when carrying out these naturalistic studies.

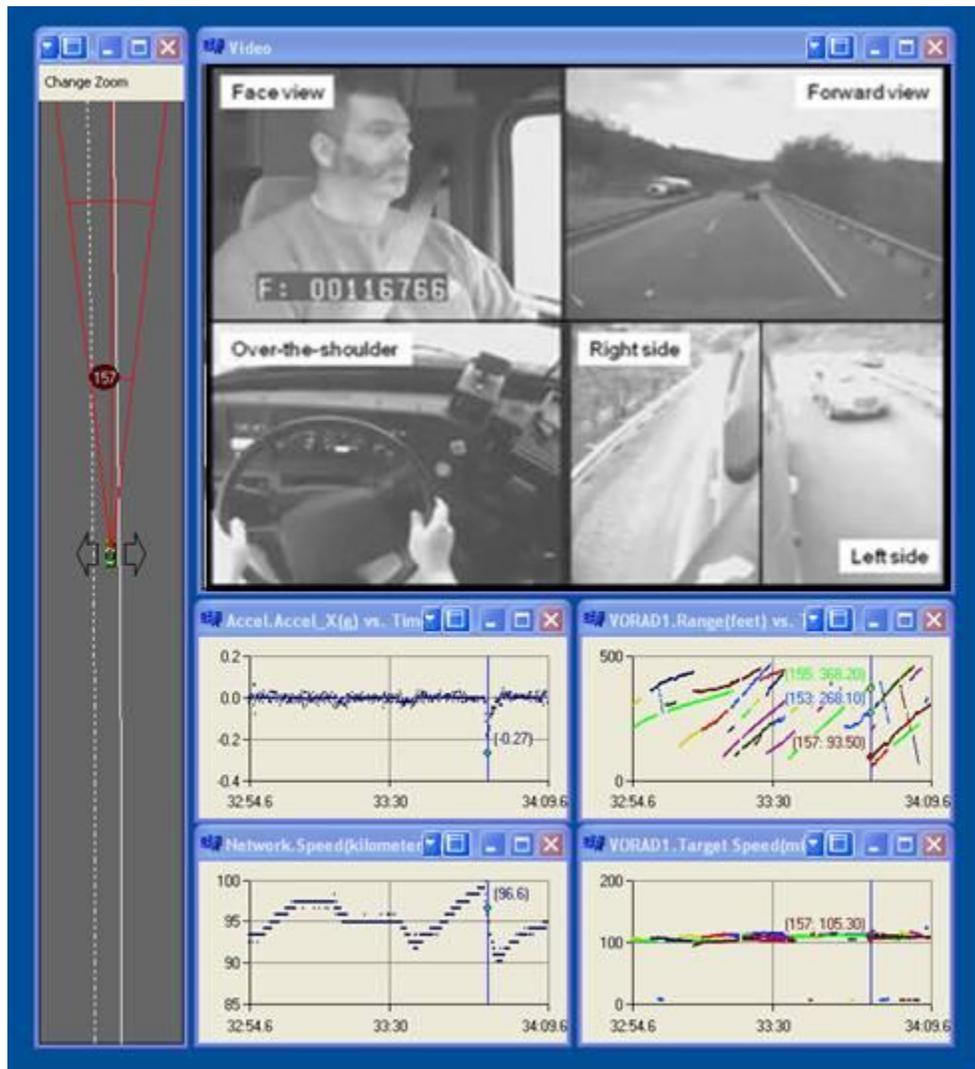


Figure 1. Screen shot. Naturalistic driving data collection.

Due to the tremendous amount of information collected during these procedures, naturalistic data collection methods require significant post-collection processing. Typically, the first step in the

data reduction process is to identify events of interest, including crashes, near-crashes, and crash-relevant conflicts. To find events of interest, VTTI developed a Data Analysis and Reduction Tool (DART) that scans the dataset using user-defined threshold algorithms to identify notable actions (e.g., hard braking, quick steering maneuvers, short times-to-collision (TTC), and lane deviations (including median encroachments). All detected events are analyzed based on “instant replays” of video data and associated dynamic data recordings of the events. This analysis captures both the observable causal sequences leading to events as well as the conditions and correlates of event occurrence.

Three datasets were used in this study:

- 8-truck study database using 100 commercial motor vehicle (CMV) drivers from 4 fleet companies, yielding 735,000 miles (14,500 hours) of driving data.
- 34-truck study database using 103 CMV drivers from 3 fleet companies, yielding approximately 2,300,000 miles of driving data.
- 100-car study database, with approximately 2,000,000 vehicle miles recorded. This data collected included 15 police-reported crashes, 67 non-police-reported crashes, and 761 near-crashes.

Nine truck drivers and eleven car drivers were selected for analysis and agent implementation based on the type of crashes and near-crashes they experienced. The drivers with the most crashes and near-crashes that were forward conflicts were selected.

EVALUATION AND EXTENSION OF EXISTING MODELS

The research compared the performance of existing car-following models to establish a baseline for comparison with the developed agents. In addition, it was envisioned that the developed agents might be used in conjunction with existing models in a hybrid simulation approach, where the agents might be invoked to simulate only safety-critical events. Special attention was devoted to the Wiedemann model because the prototype development was implemented in the VISSIM simulation package with an underlying Wiedemann model, making the Wiedemann model a good illustrative example for this effort.

Analysis of naturalistic data offered new insights into driver behavior by analyzing an exorbitant amount of car-following data for each driver. The research team discovered and proposed new thresholds for the Wiedemann model: a passing threshold and a hook following threshold. The addition of the passing threshold provides the model with the ability to easily transition from car-following to lane-changing behavior. This threshold also provides a way to force car-following behavior when a lane change is not possible. The addition of these new thresholds and the driver-specific equations improves upon the Wiedemann model’s ability to represent real-world driving behavior.

The Wiedemann model was further improved by replacing its acceleration equations in the approaching, closely approaching, acceleration following, and deceleration following regimes. The Gazis-Herman-Rothery (GHR) model was given a different set of calibration parameters for each of those regimes. The reaction time, T , was used as a calibration parameter to obtain a measure of the attentiveness of the different drivers. The combination of the Wiedemann model and the GHR model presents advantages when calibrating to the car-following behaviors of individual drivers. The hybrid Wiedemann-GHR model calibrated to four individual drivers’ results with 5 percent to 43 percent less error than the Wiedemann model alone (Figure 2).

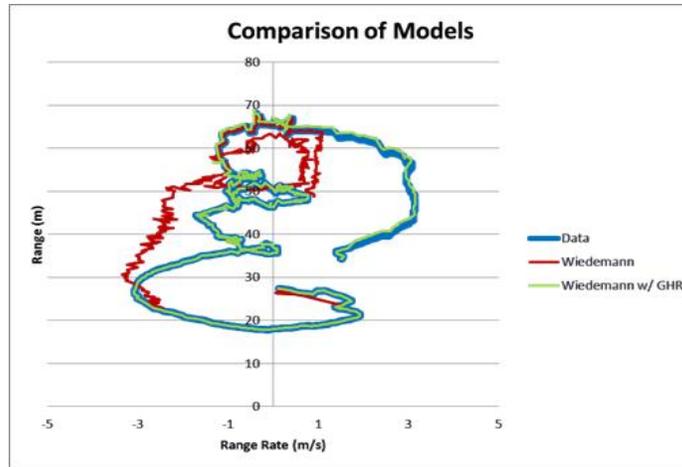


Figure 2. Graph. Comparison of car-following models.

The developed model was compared with selected, commonly used car-following models, including the GHR, Wiedemann, Fritzsche, Gipps, Intelligent Driver Model (IDM), and the velocity difference model (VDIFF). Each model was calibrated by the use of a genetic algorithm. A genetic algorithm was used because of its ability to adequately and accurately find the optimal solution when multiple parameters are present, as in some of the models. Comparison results found that different models have the least error for different drivers, but the VDIFF and the Wiedemann model both performed adequately across all the drivers tested.

The hybrid model developed in this research allows better modeling of normal car-following behavior than the comparison models. It also provides better and more seamless integration between car-following models of normal events as well as agent-based modeling of safety-critical events. The behavior of drivers during safety-critical events differs from their normal behavior and is better handled with the techniques developed in this project because, unlike the current models, the hybrid model has the capacity to address the complex dynamics that occur during a safety-critical event.

AGENT DEVELOPMENT

The research conducted in this project can be considered a foundation for agent-based modeling and simulation of driver behavior based on naturalistic data. One of the major contributions of this research is the developed integrated framework for safety and operation analysis. Agent-based modeling (ABM) is a new paradigm that consists of describing a system from the perspective of its constituent units. Compared to “top-down” modeling, agent-based modeling is bottom-up, with systems that are characterized by many agents or decision makers with dispersed control, many organizational levels, the ability of agents to adapt, and the use of internal models to anticipate the future. Agent-based modeling uses rules or equations to describe individual behaviors, so an agent-based model actually starts with a set of rules and uses those rules to generate data that can be analyzed.

The team developed agents that encapsulated individual driving behaviors of 11 car drivers and 9 truck drivers. They also developed two “mega” agents that encapsulated the behaviors of all car and truck drivers, respectively. Figure 3 illustrates the learning process for the agents.

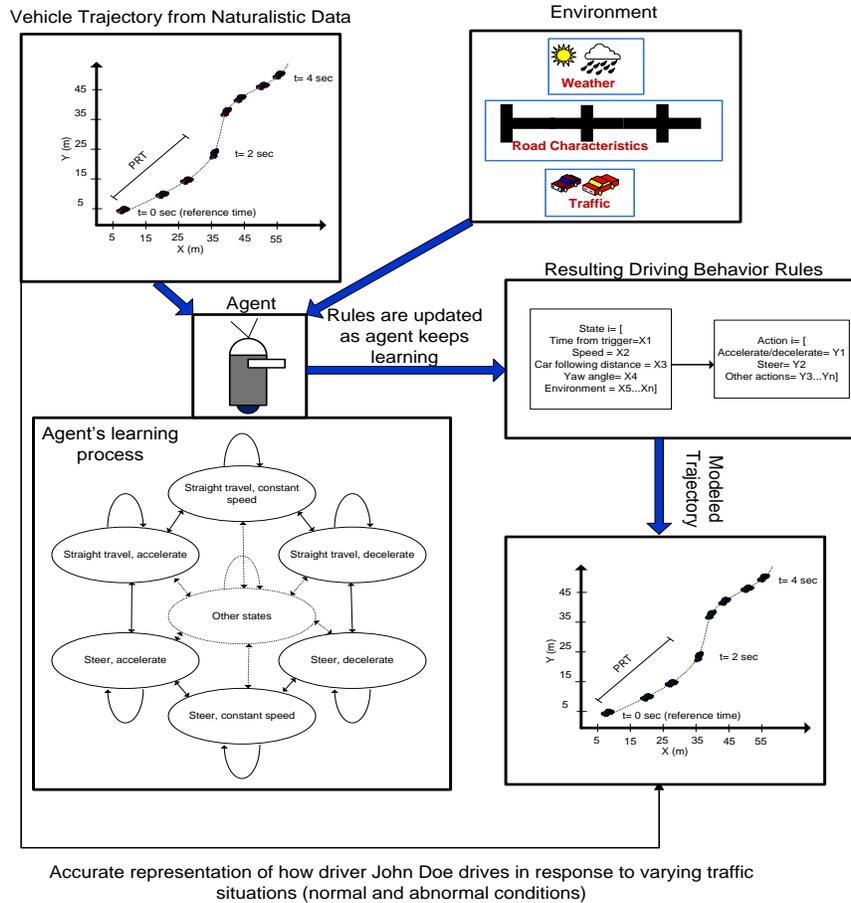


Figure 3. Illustration. Agent learning process.

The basic training for car-following behavior was handled via a back-propagation artificial neural network (BP-ANN). BP-ANN is applied to associate the nonlinear relationship between input states/actions and output states/actions. In this case, a BP-ANN network was used to determine driver longitudinal actions in normal car-following driving conditions and to estimate state transition based on current traffic state and an agent's longitudinal action. Basic training was done by extracting car-following epochs from the naturalistic driving database. The research team used both 10-Hz and 1-Hz data from one car-following episode, the same dataset that was used for the GHR car-following model calibration. The training performance for both the 10-Hz and 1-Hz datasets demonstrated stronger performance by the agent than by the GHR model, as did the model validation using a different dataset.

The research team also trained agents using neuro-fuzzy actor-critic reinforcement learning (NFACRL). The proposed and revised NFACRL approach alleviates most of the computation burden that was associated with the previously developed approaches. Traffic state variables (e.g., speed) and driver actions (e.g., acceleration) are continuous, and conventional reinforced-learning methods cannot solve high-dimensional, continuous state problems when agent actions are not drawn from a discrete set. NFACRL is able to translate high-traffic state input variables into discrete fuzzy sets and generate continuous action using a weighted average of discrete actions.

NFACRL consist of an actor-critic reinforcement learning mechanism that performs a reinforcement learning training through two parts: (1) an actor that is responsible for producing an action corresponding to each state and (2) a critic that is responsible for calculating the long-term reward for the produced action and use the reward to reinforce good actions. The *neuro* part of the NFACRL mechanism uses a neural network approach to update all the weights responsible for mapping states to actions based on the actor’s calculated reward. The *fuzzy* part of the NFACRL is responsible for transforming the continuous state space into single values representing degrees of membership in predefined discrete fuzzy sets, allowing the algorithm to handle high-dimensional state space. All the safety-critical events available from the naturalistic database were used to train 20 agents. The safety-critical events selected in this study are rear-end crash and near-crash conflict, especially when a following vehicle reacts to sudden braking or a sudden merge of the leading vehicle from the adjacent lane. When an event happens, the driver of the following vehicle brakes or swerves to avoid the incoming conflict. Using naturalistic traffic states and driving actions during crash and near-crash events, this approach was able to reproduce actual driver behavior during safety-critical events with R-squared values as high as 0.98. Cross-validation was performed by applying the training for one agent (driver) to another driver’s situation (Figure 4 and Figure 5). The performance of the 20 agents along with the mega agents is shown in Table 27 and Table 28.

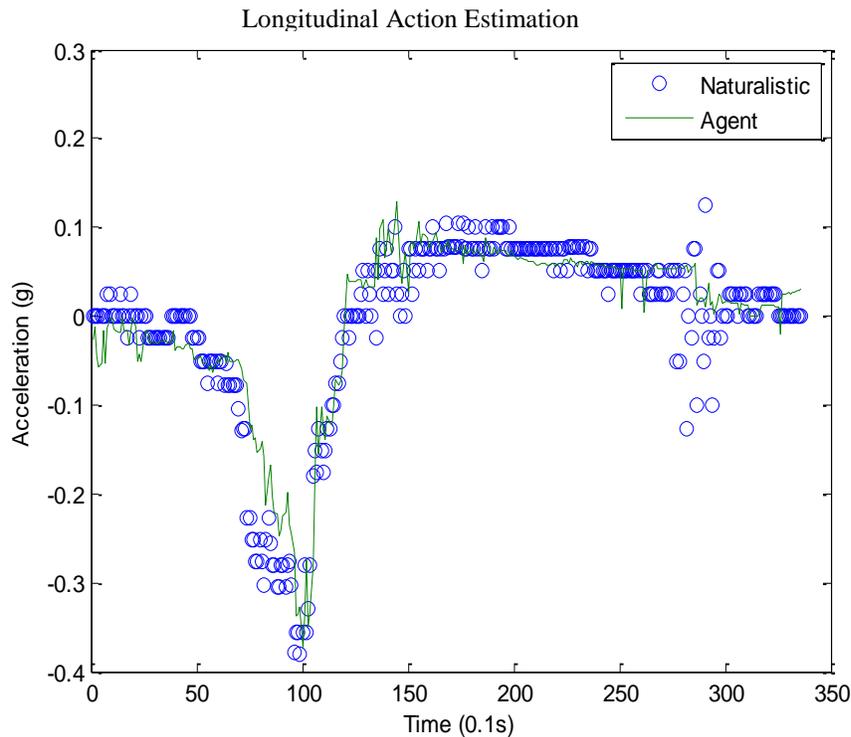


Figure 4. Graph. Acceleration of mega-agent.

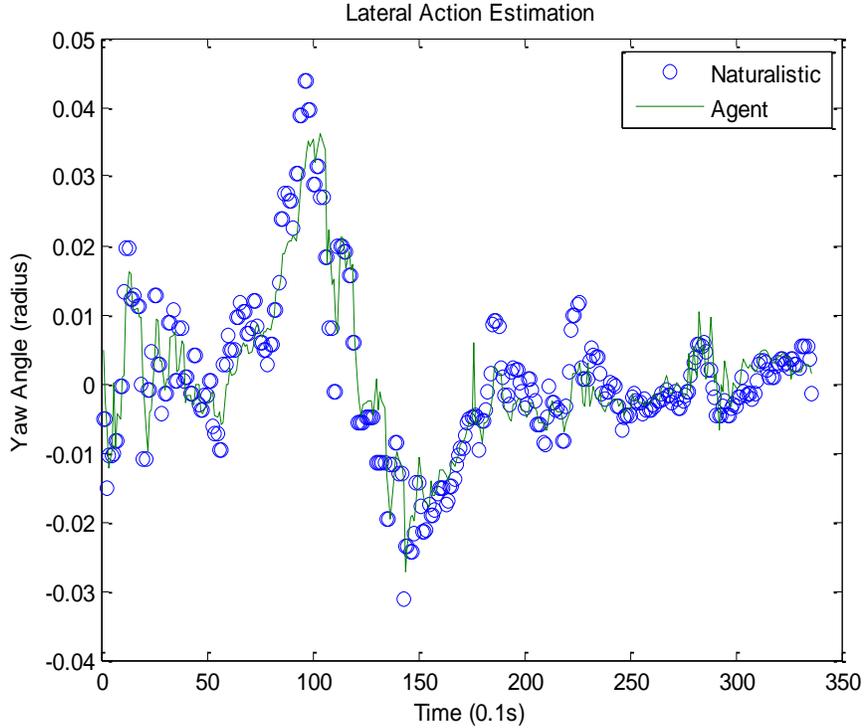


Figure 5. Chart. Yaw angle of mega-agent.

Table 1. R-squared values of the truck mega-agent.

Event	Agent Longitudinal Acceleration	Agent Lateral Acceleration	Mega-Agent Longitudinal Acceleration	Mega-Agent Lateral Acceleration	Cross Validation with Agent A Longitudinal Acceleration	Cross Validation with Agent A Lateral Acceleration
Driver A	0.97	0.97	0.98	0.97	0.98	0.97
Driver B	0.97	0.94	0.97	0.91	0.82	0.60
Driver C	0.98	0.96	0.97	0.96	0.93	0.86
Driver D	0.99	0.92	0.99	0.88	0.86	0.64
Driver E	0.88	0.96	0.81	0.95	0.47	0.76
Driver F	0.98	0.96	0.94	0.96	0.83	0.43
Driver G	0.86	0.98	0.84	0.88	0.86	0.62
Driver H	0.96	0.99	0.95	0.99	0.63	0.98
Driver I	0.95	0.98	0.93	0.98	0.48	0.75
Driver J	0.85	0.98	0.85	0.97	0.66	0.32

Table 2. R-squared values of the car mega-agent.

Event	Agent Longitudinal Acceleration	Agent Lateral Acceleration	Mega-Agent Longitudinal Acceleration	Mega-Agent Lateral Acceleration	Cross Validation with Agent D Longitudinal Acceleration	Cross Validation with Agent D Lateral Acceleration
Driver a	0.94	0.94	0.93	0.70	0.90	0.50

Driver b	0.97	0.94	0.96	0.79	0.82	0.30
Driver c	0.93	0.97	0.97	0.83	0.93	0.77
Driver d	0.97	0.93	0.96	0.93	0.97	0.93
Driver e	0.95	0.98	0.91	0.98	0.68	0.95
Driver f	0.97	0.92	0.96	0.89	0.92	0.86
Driver g	0.98	0.93	0.99	0.88	0.98	0.79
Driver h	0.98	0.93	0.99	0.89	0.98	0.90
Driver i	0.97	0.95	0.98	0.91	0.90	0.85
Driver j	0.95	0.92	0.99	0.89	0.98	0.65

The research also used a robust agent activation methodology based on the correct recognition of an impending safety-critical event using discriminant analysis. A discriminant analysis was used to find the specific combination of six variables that result in a significant difference between normal and safety-critical driving episodes that were extracted from the naturalistic database. In the examples studied, there was a visible transition from safe or normal behavior to safety-critical behavior. This resulted in the ability to turn the agent behavior on and off. The trained agents offer benefits to behaviors in safety-critical events. Furthermore, because simulation packages are designed to follow specific car-following models, the ability to turn the agent behavior on and off can be beneficial. This also means that the agents could be trained only for the safety-critical event behavior, and the car-following models can be used otherwise, which makes the application of the agents more flexible, bridging the gap between the current state of practice and the new trained agent method.

Finally, the research developed the agent-based simulation components integrated with the VISSIM simulation package through its driver model dll. Parameter files for 20 agents were developed and delivered as part of this project. Evaluation of the developed agent prototypes showed that the agents would adopt more realistic decelerations than the Wiedemann model.

SIMULATION AND ANALYSIS

Developed agents were implemented in VISSIM and compared with the existing VISSIM car-following model. Evaluation metrics included the following:

- Minimum following distance (m).
- Maximum deceleration (m/s^2).
- Average acceleration (m/s^2).
- Average of absolute acceleration (m/s^2).
- Average speed (km/hr).

The Wiedemann car-following model did not produce any conflicts when the time-to-collision threshold was set to 1.5 s. The agent-based models produced conflicts, with most of the agents occasionally entering the “emergency” regime, where the agent’s speed is greater than that of the lead vehicle, but the following distance is small. The agent’s deceleration rates were realistic and have resulted in several collision and run-off-the-road incidents.

Summary of Operations Analysis

The qualitative and quantitative analysis on the simple linear VISSIM network revealed the following findings:

- Most agents were able to enter the emergency regime, where an agent's speed is greater than the lead vehicle's yet headway is small. Vehicles following the Wiedemann model as implemented in VISSIM were unable to enter the emergency regime, and instead vehicles were forced to decelerate at unrealistic rates to avoid collisions.
- Almost all agents were able to eliminate the improbably high deceleration rates used by vehicles following the Wiedemann model as implemented in VISSIM. The agents instead ran over, swerved, or collided with leading vehicles. The vast majority of the agents' deceleration rates to avoid collisions were realistic.
- Although agents experienced realistic deceleration rates, in many cases they did not behave consistently across all simulation runs. The agents' behavior indicates that although each agent is programmed to respond a certain way to a certain scenario, the definition of that scenario may be sufficiently narrow that changes in the random seed fail to re-create it. The definition of the scenario may need to be broadened or better defined so that vehicles can behave in a more consistent way. This issue can be addressed by exposing the agents to more data points during the training process.
- The agents were trained to emphasize safety-critical modeling, and therefore the frequency of crashes and near crashes during an agent-based episode was exaggerated. This was intentional in this project to showcase the ability of agents to simulate safety-critical events. However, for the purpose of real analysis, the agents should take into account the realistic probability of crashes and produce appropriate statistics. This should be a focus for future research.

Summary of Safety Surrogate Analysis Model

The safety surrogate measure analysis on the simple linear VISSIM network revealed the following findings:

- As expected, the agents produced many more safety-critical events than when VISSIM used the Wiedemann car-following model to control the following vehicle. The Wiedemann car-following model never produced any conflicts when the TTC threshold for a conflict was set at 1.5 s, regardless of the scenario being evaluated.
- Several of the agents were not well suited to analysis using Surrogate Safety Assessment Model (SSAM). Agents that departed the roadway by a significant margin were not properly classified as having caused safety-critical events.
- When the TTC threshold for conflicts was increased to 2.5 s, the Wiedemann model sometimes produced conflicts. However, the number and severity of conflicts with the Wiedemann model were not as severe as when the agent was active.
- Reactions of the agents following the safety-critical event were sometimes not realistic. Agents would repeat behaviors that created safety-critical events multiple times, or would proceed following a collision as if nothing had happened. Before the agents can be implemented in a more complex model, behaviors following safety-critical events will need to be re-examined for the agents. In effect, noncollision events would need to reset

to the traditional Wiedemann model, and collision events would need to be modeled by having both vehicles come to a stop after the event.

The team suggests that before the agents are ready to be implemented in a production environment in VISSIM, several issues need to be resolved:

1. The areas of application (speeds, geometric conditions) where specific agents should be used need to be defined. Before being used in a production environment, the selection of appropriate agents for a specific condition should be automated.
2. As a future research item, there is a need to estimate the frequency with which certain agent behaviors occur in the driving population to ensure that they are modeled at appropriate levels.

Once these changes are made, the inclusion of agents in VISSIM would offer a powerful opportunity to examine safety implications of different designs, as well as to better assess the impacts of incidents and other unpredictable driver behavior on travel time reliability.

CONCLUSIONS AND RECOMMENDATIONS

There are several suggestions for future research that can be categorized under theoretical, applied, and related efforts. On the theoretical aspect, the next step of this research is the extension of the developed NFACRL framework to simulate other traffic behavior, such as lane-changing behavior and merging behavior in the upstream and downstream of ramps.

On the applied aspect, our initial evaluations on events beyond what the agents were trained for revealed that agents can behave erratically in those cases. It is therefore recommended to train agents using more detailed naturalistic datasets (e.g., Strategic Highway Research Program 2 naturalistic datasets). More importantly, research should be conducted to evaluate the potential of using agent-based simulation to examine different design alternatives based on safety and operation factors. The evaluation effort should examine unique or problematic geometry, network design characteristics, and traffic conditions that necessitate both operational and safety analysis (e.g., different weather conditions) and compare the results to existing and/or known outcomes.

Finally, on the related research aspect, the research conducted in this report has shown the importance of two key future research issues:

1. Adaptability of agents in real time: agents could be programmed to adapt during the simulation itself.
2. Human factor issues related to warning individual drivers about a change in their driving behavior that might lead to a safety-critical event.

Future recommendations include the following:

- Assess the frequency that certain agent behaviors occur in the actual driving population.
- Extend the NFACRL framework to simulate other traffic behavior such as lane changing and merging.
- Examine the adaptability of agents in real time, i.e., reprogram them during the simulation.
- Examine human factor issues related to warning individual human drivers about a change in their driving behavior that might lead to a safety-critical event.

CHAPTER 1. INTRODUCTION

BACKGROUND

Existing traffic analysis and management tools do not model the ability of drivers to recognize their environment and respond to it with behaviors that vary according to the encountered driving situation. The small body of literature on characterizing drivers' behavior is typically limited to specific locations (i.e., by collecting data on specific intersections or freeway sections) and is very narrow in scope. The majority of traffic modeling and parameter calibration research has assumed somewhat similar driving conditions and behavioral sets for the entire population of drivers. Differences in drivers' actions are represented merely by drawing samples from statistical distributions assigned to each driver type. The number and properties of driver types used in this approach becomes a calibration issue, which might or might not improve the modeling effort. This approach does not capture or predict the actions of individual drivers responding differentially to situational factors.

A comprehensive representation of driver behavior varying in response to encountered driving situations is not available in the literature and, therefore, reproduction of such behaviors using statistical parameters is not adequate. At any instant, a driver's action will be influenced by the driver's perception of the surrounding environment or changes in the driving situation, the vehicle's position along its driving path, and the vehicle's position relative to other vehicles.

This report documents the research findings of the Federal Highway Administration Exploratory Advanced Research Program's "Driver Behavior in Traffic" study. The objective of this research was to characterize driver behavior under naturalistic driving conditions. More specifically, the research aimed to determine critical parameters related to freeway driving during critical events, including target speed selection, accelerations, lane changing behavior, car-following distance, response times, and emergency stopping behavior for representative conditions. Of special interest to this research was the integrated modeling of both normal and safety-critical driving. The techniques used in this research, namely agent-based simulation techniques, allow modeling of individual driver behavior beyond existing models' capabilities.

OVERALL SCOPE OF THE EFFORT

The scope of the work was to develop and implement intelligent agents that can encapsulate individual drivers' decisions in response to varying traffic situations. The developed agents were designed to learn individual drivers' actions for any given traffic state retrieved from the Virginia Tech Transportation Institute's (VTTI) database of naturalistic driving data. These characteristic driving rules of the agents were coded in a VISSIM simulation environment to test and study their effects on agents' behavior in different conditions and scenarios (as illustrated in Figure 6). The goal of this effort is to provide the industry with methods for developing more accurate and more sensitive traffic simulation models. This could also enable future research to develop new generations of traffic simulation models that accurately model driver behavior during incidents and other complex traffic situations.

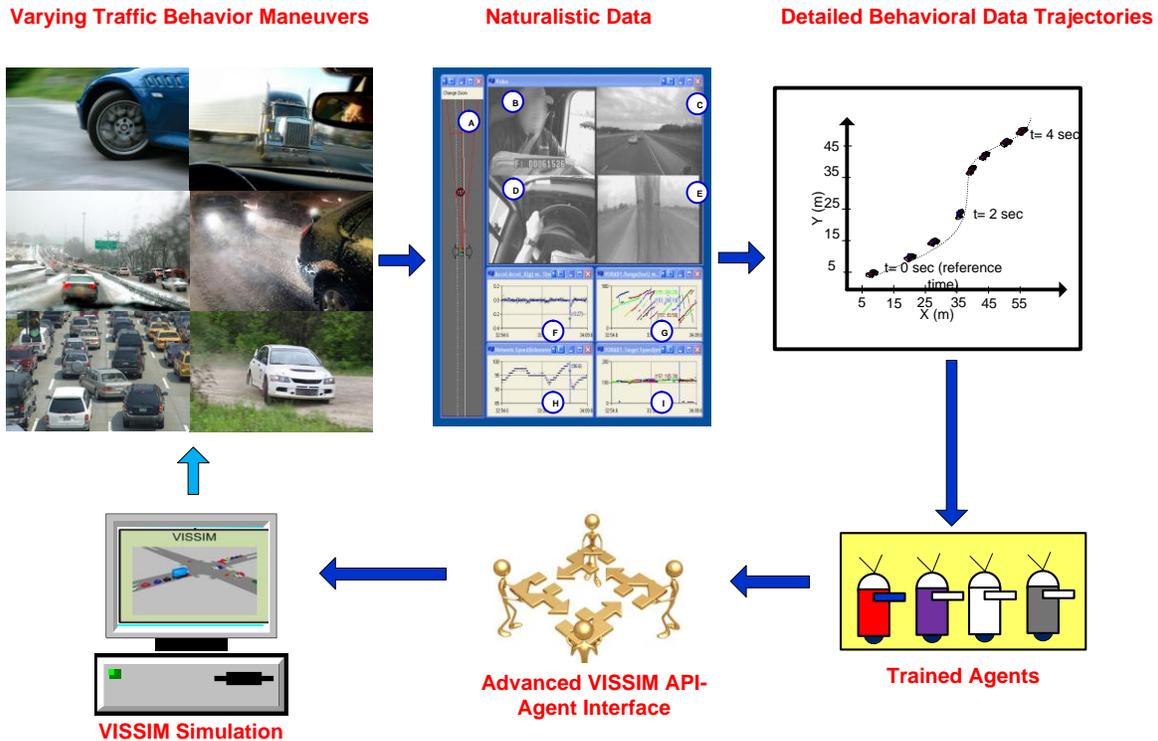


Figure 6. Illustration. System representation of driver's response to encountered traffic situation.

REPORT ORGANIZATION

This report is organized as follows: Chapter 2 provides a synthesis of the literature on key domestic and international past and continuing studies on car-following and lane-changing microscopic models, naturalistic data collection efforts, driver behavior and characteristics, artificial intelligence learning techniques, and agent-based simulation. Chapter 3 describes the VTTI naturalistic driving data processing and provides a summary of the car and truck drivers that were selected for further analysis. Chapter 4 provides a detailed description of the Wiedemann car-following model and outlines the procedure used to calibrate the model. The chapter also describes how an improved Wiedemann model was developed based on the information found in the naturalistic driving data. Chapter 5 describes the development of intelligent agents that emulate drivers' actions for any given traffic state. Artificial intelligence techniques, such as reinforcement learning (RL), artificial neural networks, and fuzzy logic implementations are also described. Chapter 6 discusses robust statistical techniques to activate the developed agents for simulating safety-critical events. The chapter describes the use of discriminant analysis to find the difference between normal driving behavior and safety-critical event behavior. Chapter 7 describes the agent implementation in the VISSIM platform to validate the research methodology and identify potential implementation issues for future research. Chapter 8 describes the evaluation of the developed prototypes and compares the results of agent-based simulation with Wiedemann (normal) simulation. Finally, Chapter 9 provides a summary of the findings of the research and recommendations for future work.

CHAPTER 2. LITERATURE REVIEW

INTRODUCTION

Existing traffic analysis and management tools do not model drivers' ability to recognize their environment and respond to it differently depending on the situation. There is very limited literature on characterizing driver behavior. What research exists is typically limited to specific locations (i.e., by collecting data on specific intersections or freeway sections) and is very limited in scope. The majority of traffic modeling and parameter calibration research has assumed somewhat similar driving conditions and behavioral sets for the entire driver population. Differences in drivers' actions are represented merely by drawing samples from statistical distributions assigned to each driver type. The number and properties of driver types used influences calibration, and this influence might or might not improve the modeling effort. This approach does not capture or predict the effects of situational factors on individual drivers' actions. A comprehensive theoretical representation of drivers' varying behavior in response to encountered driving situations is very limited and, therefore, reproduction of such behaviors (or actions) using statistical parameters is not adequate.

This chapter synthesizes a literature review conducted by the research team on key domestic and international studies on car-following and lane-changing microscopic models, naturalistic data collection efforts, driver behavior and characteristics, and artificial intelligence learning techniques, especially reinforcement learning.

CAR-FOLLOWING MODELS

In the last 50 years, a considerable amount of research has focused on modeling longitudinal driver behavior, producing a large number of car-following models⁽¹⁾, including Gazis-Herman-Rothery (GHR) models, safety distance models, linear models, and psychophysical or action point models. Most microscopic models assume that human drivers react to the stimuli from neighboring vehicles with the dominant influence originating from the directly leading vehicle. This is known as the "follow-the-leader" or "car-following" approximation. Other models introduced "safe time headway" and a desired velocity. Calibrating these car-following models required different levels of effort, and the results depend on data availability, calibration method, and model structure. The GHR model, for instance, requires the calibration of two parameters, limiting the application of the model. The Gipps model⁽²⁾ (a safe-distance model) uses two different transfer functions for reproducing free-flow and car-following conditions. Psychophysical models, on the other hand, assume that the driver will perform an action when a threshold (a function of speed difference and distance) is reached. Estimating distances and speed differences among drivers makes it difficult to calibrate the individual threshold associated with this model. Linear regression and heuristic algorithms (e.g., genetic algorithms) are two widely used methods for model calibration. Despite different mechanisms and software interfaces, when multiple simulation software applications were compared, the resulting behaviors of the models showed similarities.⁽³⁾ In addition, it seems that error cannot be eliminated even if more parameters are introduced.

The importance of not only developing calibration procedures but also interpreting obtained calibration results is discussed in Ossen.⁽⁴⁾ The three main findings were “(1) measurement errors can yield a considerable bias in the estimation results; (2) parameters minimizing the objective function do not necessarily capture following dynamics best; (3) measurement errors substantially reduce the sensitivity of the objective function and consequently reduce the reliability of the results.”

A significant part of the deviations between measured and simulated trajectories can be attributed to the interdriver variability and the intradriver variability (human drivers do not drive constantly over time, and their behavioral driving parameters change). The latter accounts for a large part of the deviations between simulations and empirical observations.⁽⁵⁾

The first car-following model describing car following in terms of vehicle trajectory rather than just velocities and spacing was proposed by Herrey and Herrey in 1945.^(6; 7) Without extensive experimental basis, they proposed that a driver must maintain a “safe driving distance” (including the stopping distance), which results in a quadratic relationship between spacing and speed. The safety distance models, or collision avoidance model, represents a model in which the driver of the following vehicle always keep a safe distance to the vehicle in front.

Pipes Model

In 1953, Pipes⁽⁸⁾ developed one of the first car-following models using the California Motor Vehicle Code rules: “allow at least the length of a vehicle between the subject and leading vehicle for every ten miles per hour speed at which the subject vehicle is traveling” assuming a vehicle length of 20 feet. The model is divided into two equations. The first equation applies when the lead vehicle has a constant acceleration while trying to reach a cruising speed. The second equation applies after the lead vehicle reaches and maintains the cruising speed. Both equations estimate the acceleration of the following vehicle at time t based on the cruising speed and acceleration time of the lead vehicle. This model is presented in a fashion that is applicable to a line of traffic with each vehicle accelerating according to the action of the vehicle in front of it.

$$a_{k+1}(t) = \frac{v_m}{T_0} G_k \left(\frac{t}{T} \right) \text{ if } 0 < t < T_0$$

Equation 1.

$$a_{k+1}(t) = \frac{v_m}{T_0} \left[G_k \left(\frac{t}{T} \right) - G_k \left(\frac{t - t_0}{T} \right) \right] \text{ if } t > T_0$$

Equation 2.

where

t = time,

T = a time constant ≈ 1 ,

a_{k+1} = acceleration of the $k+1$ (following) vehicle,

v_m = constant velocity of lead vehicle (cruising speed),

T_0 = time for lead vehicle to reach cruising speed at a constant acceleration from rest at $t=0$, and

$$G_k(t) = 1 - e^{-t} \left[1 + \frac{t}{1!} + \frac{t^2}{2!} + \dots + \frac{t^{(k-1)}}{(k-1)!} \right]$$

Equation 3.

where

$G_k(t)$ is the ratio of an incomplete to a complete gamma transform function.

Most of the early work in car following assumed that the driver is able to consider the space headway and the relative speed between his car and the lead vehicle. To model this stimulus-response model, several approaches have been developed, including linear, proportional approaches and rule-based approaches based on fuzzy logic. In general, these car-following models compute the reaction of a following vehicle to the actions or stimulus of the leading vehicle as Response (t) = Sensitivity X Stimulus (t-T). The response of the following vehicle is translated into an acceleration or deceleration in proportion to the stimulus after a time lag T (reaction time of the driver of the following vehicle).

The GM models

In 1958, a series of models were developed at the GM Research Laboratory by Chandler, Herman, and Montroll⁽⁹⁾; Herman and Potts⁽¹⁰⁾; Gazis, Herman, and Potts^(11; 12); and Gazis, Herman, and Rothery¹². The most important of these is the generalized linear model, usually known as the GM model. In fact, most of the other models, for example Pipes⁽⁸⁾ and Forbes⁽¹³⁾ are special cases of the GM model.

Chandler, Herman, and Montroll⁽⁹⁾ proposed a linear model that assumes that the acceleration of the following car depends exclusively on the speed difference between the cars.

$$a_n(t) = \alpha \Delta v_n^{front}(t - \tau_n)$$

Equation 4.

where

$a_k(t)$ = acceleration applied by the driver n at time t ,

α = constant,

$$\Delta v_n^{front}(t - \tau_n) = v_n^{front}(t - \tau_n) - v_n(t - \tau_n),$$

Equation 5.

$$v_n^{front}(t - \tau_n) = \text{subject speed at time } (t - \tau_n),$$

Equation 6.

t = time of observation, and

τ_n = reaction time for driver n .

Gazis, Herman, and Potts⁽¹⁴⁾ developed the GM 2 parameter model (GM2P), in which the linear model constant sensitivity term was further developed to consider the macroscopic speed density relationship and the space headway. The incorporation of the space headway implies that the driver is more sensitive to the action of the leading vehicle in the case of smaller headways.

$$a_n(t) = \frac{\alpha}{\Delta x_n(t-\tau_n)} \Delta v_n^{front}(t - \tau_n)$$

Equation 7.

where

$\Delta x_n(t-\tau_n)$ is the space headway at time $(t-\tau_n)$ and

α = constant, estimated for each dataset using correlation analysis.

The model of Edie, or the GM 3 parameter model (GM3P), addresses the shortcoming of the linear model of free-flow speed at zero density and assumes that the reaction of the following car depends also on the speed of the following car and not only on the speed difference and the relative distance between the two cars. The sensitivity is proportional to the speed and inversely proportional to the square of the relative distance.

$$a_n(t) = \alpha \frac{v_n(t)}{\Delta x_n(t - \tau_n)^2} \Delta v_n^{front}(t - \tau_n)$$

Equation 8.

The Gazis, Herman and Rothery model, usually referred to as the GHR model or as the general car-following model, was introduced in 1961 and is represented by the following equation.

$$a_f(t + T_r) = c v_f^m(t + T_r) \frac{\Delta v(t)}{\Delta x^l(t)}$$

Equation 9.

where

T_r = time between the observation of a certain stimulus and the reaction to that stimulus,

$a_f(t + T_r)$ = acceleration of the following vehicle at time $(t + T_r)$,

$v_f(t + T_r)$ = speed of following vehicle at time $(t + T_r)$,

$\Delta v_n(t)$ = relative speed between the following car and the car immediately in front ($v_{leader} - v_{follower}$),

$\Delta x(t)$ = relative distance between following car and car immediately in front ($x_{leader} - x_{follower}$), and

m, l, c = parameters describing the car-following behavior.

The GHR model is therefore a general form of the early models. When $l=m=0$, the GHR model represents that of Chandler, Herman and Montroll; when $m=0$ and $l=1$, the model represents Gazis, German, and Potts' model; and when $m=1$ and $l=1$, the model represents Edie's model.

The Wiedemann Model

In 1974, Wiedemann⁽¹⁵⁾ introduced a car-following model that is based on psychophysical behavior. The Wiedemann model addresses two unrealistic assumptions of the GM models: (1) the driver follows the leading vehicle whatever the spacing is and (2) the driver has perfect reaction and response time. The Wiedemann model has been improved since then by Leutzbach⁽¹⁶⁾ by introducing perceptual thresholds as minimum values of the stimulus. The concept of thresholds in the Wiedemann model captures the driver's alertness in conditions of small space headway and the lack of explicit car-following behavior in large headway. In addition, it allows the model to explain the oscillation phenomena observed in car-following behaviors. The following diagram (Figure 7) shows Wiedemann car-following thresholds with respect to vehicle speed difference and space headway.

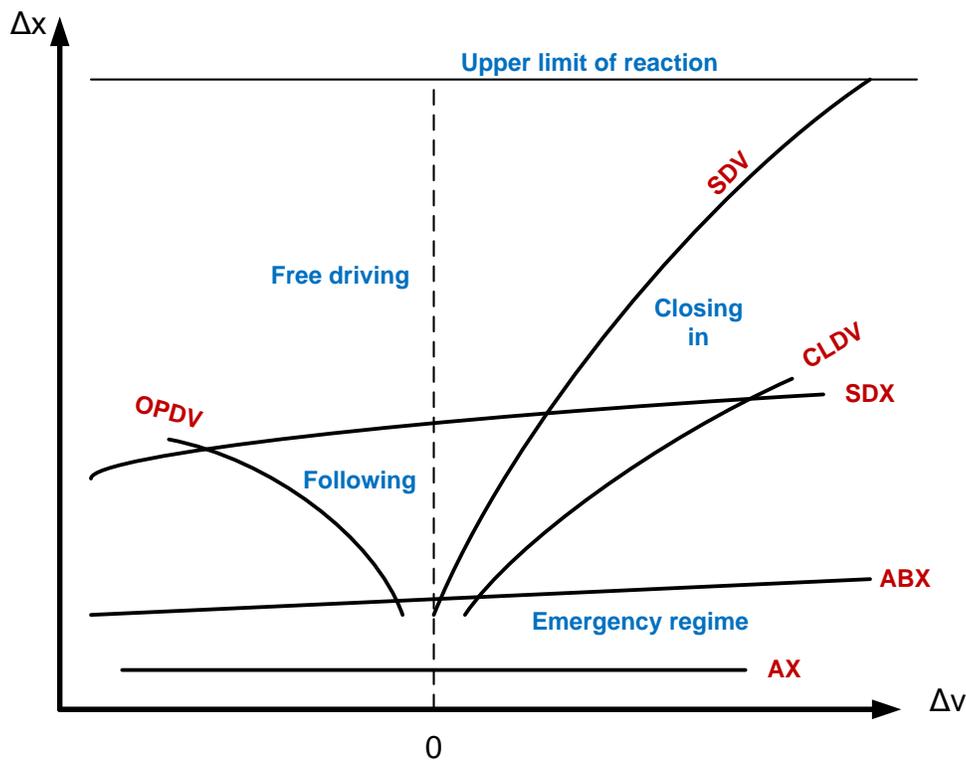


Figure 7. Graph. Wiedemann car-following model's thresholds (AX, ABX, SDX, CLDV, SDV, OPDV).⁽¹⁷⁾

The following is a definition of the Wiedemann car-following model's thresholds:

1. The threshold AX represents the desired distance between stationary vehicles. This distance consists of the leading vehicle length and the desired front-to-rear distance.

$$AX = L_{n-1} + AXadd + RND1_n \times AXmult$$

Equation 10.

where

L_{n-1} = the leading vehicle length,

AXadd and AXmult are calibration parameters, and

RND1_n is a normal distributed random number for vehicle (n).

2. The threshold ABX represents the desired minimum following distance at low speed differences.

$$ABX = AX + BX$$

Equation 11.

$$BX = (BXadd + BXmult \times RND1_n) \times \sqrt{v}$$

Equation 12.

$$v = \begin{cases} v_{n-1} & \text{for } v_n > v_{n-1} \\ v_n & \text{for } v_n \leq v_{n-1} \end{cases}$$

Equation 13.

where

BXadd and BXmult are calibration parameters and

v = vehicle speed.

3. The threshold for maximum following distance, SDX, fall between 1.5 and 2.5 times ABX

$$SDX = AX + EX \times BX$$

Equation 14.

$$EX = (EXadd + EXmult \times (NRND - RND2_n))$$

Equation 15.

where

EXadd and EXmult are calibration parameters,

NRND is a normal distributed random number, and

RND2_n is normal distributed driver dependent parameter.

4. The threshold approaching point, SDV, represents the point where the driver notices that he/she is approaching a slower vehicle.

$$SDV = \left(\frac{\Delta x - L_{n-1} - AX}{CX} \right)^2$$

Equation 16.

$$CX = CXconst \times (CXadd + CXmult \times (NRND1_n + RND2_n))$$

Equation 17.

where

CXconst, CXadd, and CXmulti are calibration parameters.

5. The decreasing speed differences (CLDV) threshold provides a mechanism for a different behavior when the driver approaches the leading vehicle more closely than the approaching point. In VISSIM, CLDV is ignored and simply assumed to be equal to SDV.
6. The increasing speed difference (OPDV) threshold describes the point where the driver observes that he/she is traveling at a lower speed than the leader.

$$OPDV = CLDV \times (-OPDVadd - OPDVmult \times NRND)$$

Equation 18.

where

OPDVadd and OPDVmult are calibration parameters and

NRDV is a normally distributed random number.

The Fritzsche Model

Fritzsche⁽¹⁸⁾ developed a similar psycho-physical model that has been used in Paramics microsimulation software. The model has multiple regimes that apply to different traffic situations (i.e., normal car-following, free flow, closing in, etc., as shown in Figure 8). Fritzsche's model accounts for human perception in the definitions of the model regimes, as drivers perceive only speed differences with certain magnitude. Fritzsche defined the thresholds for perception as: perception threshold of positive speed difference (PTP) and perception threshold of negative speed difference (PTN). The model also incorporates four thresholds for the follower's space headway.

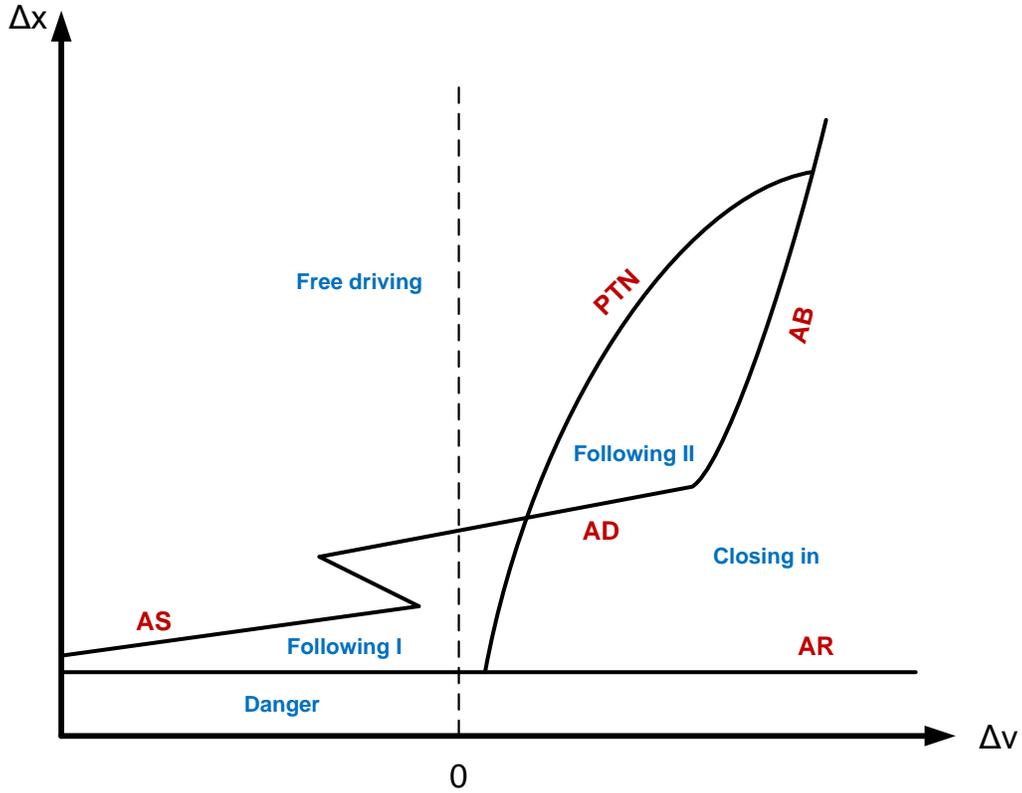


Figure 8. Graph. Different thresholds and regimes in the Fritzsche car-following model.⁽³⁾

1. Desired distance threshold, AD , represents the gap, which following driver wants to maintain.

$$AD = S_{n-1} + T_D \times v_n$$

Equation 19.

where

T_D = a parameter representing the desire time gap,

S_{n-1} = effective length of the leading vehicle, and

v_n = following vehicle speed.

2. The risk distance threshold, AR , represents the distance at which the driver applies heavy deceleration to avoid collisions.

$$AR = S_{n-1} + T_r \times v_{n-1}$$

Equation 20.

where

T_r = a parameter representing the risky time gap.

3. The safe distance threshold, AS, represents the smallest headway where positive acceleration is accepted if the distance between the vehicles is increasing.

$$AS = S_{n-1} + T_s \times v_n$$

Equation 21.

where

T_s = a parameter representing the safe distance threshold.

4. The risk braking distance threshold, AB, represents the distance where maximum deceleration is limited. Therefore collision is possible if the initial speed difference is large.

$$AB = AR + \frac{\Delta v^2}{\Delta b_m}$$

Equation 22.

where

$$\Delta b_m = |b_{min}| + a_{n-1}$$

Equation 23.

b_{min}, a_{n-1} = a parameter controlling maximum deceleration.

The CA Model

The cellular automaton model, usually know as CA model,⁽¹⁹⁾ simulates freeway traffic. Monte Carlo simulations of this model show a transition from laminar traffic flow to start-stop waves with increasing vehicle density.

The CA model can be described as following seven steps.^(19; 20)

1. Determination of the randomization probability.

$$p(v_n(t), t_h, t_s, S_{n+1}(t)) = \begin{cases} p_b: \text{if } S_{n+1} = 1 \text{ and } t_h < t_s \\ p_0: \text{if } v_n = 0 \text{ and } t_{st} \geq t_{k,c} \\ p_d: \text{in all other cases} \end{cases}$$

Equation 24.

where

$$t_h = \frac{d_n}{v_n(t)}, t_s = \min(v_n(t), h_k),$$

Equation 25.

where t_{st} denotes time vehicle stops.

2. Acceleration. Determine the speed of vehicles in the next time step. $S_n(t)$ is also taken into consideration. $S_n(t)$ is determined in Step 5.

$$\text{If } (S_{n+1}(t) = 0) \text{ or } (t_h \geq t_s) \text{ then } v_n(t + 1) = \min(v_n(t) + a_k, v_{k,max})$$

Equation 26.

$$\text{Else } v_n(t + 1) = v_n(t)$$

Equation 27.

where $v_{k,max}$ and a_k are the maximum speed and acceleration capacity, respectively, of vehicles of type k .

3. Deceleration. Set speed restriction when the vehicle in front is too close, thus locating the vehicle within the effective distance (d_n^{eff}).

$$v_n(t + 1) = \min(d_n^{eff}, v_n(t + 1))$$

Equation 28.

4. Randomization.

$$\text{If } (rand() < p) \text{ then } v_n(t + 1) = \max(v_n(t + 1) - 1, 0)$$

Equation 29.

5. Determination of vehicle status identifier $S_n(t)$ in next time step.

$$S_n(t + 1) = \begin{cases} 0 & \text{if } v_n(t + 1) > v_n(t) \\ S_n(t) & \text{if } v_n(t + 1) = v_n(t) \\ 1 & \text{if } v_n(t + 1) < v_n(t) \end{cases}$$

Equation 30.

6. Determination of time stuck inside the jam (t_{st}).

$$t_{st} = \begin{cases} t_{st} = t_{st} + 1 & \text{if } v_n(t + 1) = 0 \\ t_{st} = 0 & \text{if } v_n(t + 1) > 0 \end{cases}$$

Equation 31.

7. Update position.

$$x_n(t + 1) = x_n(t) + v_n(t + 1)$$

Equation 32.

where

n = n th vehicle,

$n+1$ = vehicle in front,

k = type of vehicle,

P_b = accounting for impact of decelerating vehicle in near front,

P_0 = reflect the delay-to-start behaviors of vehicles stuck in traffic jam,

P_d = other situations,

P_c = probability of lane change,

$t_h = d_n/v_n(t)$ = time headway of n th vehicle to front,

h_k = preset time threshold of vehicle of type k in reflecting the effect of synchronized distance,

$t_s = \min(v_n(t), h_k)$ = final time threshold for initiating the consideration of the front brake light effect, taking the vehicular speed into consideration,

t_{st} = accumulated time of vehicle stuck in traffic jam,

$t_{k,c}$ = time threshold of vehicle of type k for initiating the stop-to-start behavior,

$\tau^{(safe)}$ = safe time gap for collision prevention,

Δt = duration of an individual time step,

d = space headway,

d_n^{eff} = effective distance of n th vehicle,

g = distance gap,

x = position of n th vehicle,

\dot{x} = speed (first derivative of vehicular position),

\ddot{x} = acceleration/deceleration (second derivative of vehicular position), and

Δx = basic unit of roadway length.

The SK Model

The safety condition in the CA car-following model is imposed by the term $s_{gap}(t)$,⁽²¹⁾ setting the maximum velocity equal to the space gap between two successive vehicles regardless of the required deceleration. To address the fact that the CA model fails to reproduce metastable states of very high throughput and hysteresis effects existing in traffic flow, the model was generalized to provide a one-parametric family of models with a part that reproduces metastable states and hysteresis usually known as SK Model⁽²²⁾ and represented by the following:

$$v \leq v_{safe} = b(\alpha_{safe} + \beta_{safe})$$

Equation 33.

$$\alpha_{safe} = \left[\sqrt{2 \frac{d_p + g}{b} + \frac{1}{4}} - \frac{1}{2} \right]$$

Equation 34.

$$d_p = b(\alpha_p \beta_p + \frac{\alpha_p(\alpha_p - 1)}{2})$$

Equation 35.

where

b = maximum deceleration

d_p = braking distance of lead car

α_p, β_p = integer and fractional part, respectively, of v_p/b , and

$\alpha_{safe}, \beta_{safe}$ = integer and fractional part of

$$\alpha_{safe} = \left[\sqrt{2 \frac{d_p + g}{b} + \frac{1}{4}} - \frac{1}{2} \right]$$

Equation 36.

The SK model is a development of the CA model⁽²²⁾. In place of the CA model's Step 7, the SK model uses the updating rules below:

$$v_1 = \min[v(t) + b, v_{max}, v_{safe}]$$

Equation 37.

$$v_0 = v_1 - \epsilon\{v_1 - [v(t) - b]\}$$

Equation 38.

$$v(t + 1) = v_{ran, v_0, v_1}$$

Equation 39.

$$x(t + 1) = x(t) + v(t + 1)$$

Equation 40.

where

ε = parameter,

v_1 = optimal velocity for next update,

v_{\max} = maximum velocity of cars,

v_{safe} = maximum safe velocity, and

v_{ran}, v_0, v_2 = a random number between v_0 and v_1 = maximum decel.

The IDM Model

The Intelligent Driver Model (IDM)⁽²³⁾ is defined by the following acceleration function:

$$a_n(t + \tau) = a \left[1 - \left(\frac{v_n(t)}{v^*} \right)^\delta - \left(\frac{\Delta x_{\min}^*(v_n(t), \Delta v_{n, n-1}(t))}{\Delta x_{n-1, n}(t)} \right)^2 \right]$$

Equation 41.

$$\Delta x_{\min}^*(v_n(t), \Delta v_{n, n-1}(t)) = d + T_{\text{safe}} * v_n(t) + \frac{v_n(t) \Delta v_{n, n-1}(t)}{2\sqrt{a^{\max} x_* b^{\max \text{ abs}}}}$$

Equation 42.

where

$\Delta v_{n, n-1}(t)$ = approaching rate of the following vehicle,

v^* = desired speed,

T_{safe} = safe time headway,

a^{\max} = maximum desired

$b^{\max \text{ abs}}$ = absolute maximum desired deceleration of following vehicle,

δ = acceleration component,

$\Delta x_{n-1, n}(t)$ = distance headway, and

d = vehicle length.

This expression combines the acceleration strategy in Equation 43 toward a desired speed v^* on a free road, with the parameter a for the maximum acceleration. It is further combined with the

braking strategy in Equation 44, which is dominant if the current gap $\Delta x_{n-1, n}(t)$ becomes smaller than the desired minimum gap.

$$\dot{a}_{free} = a \left[1 - \left(\frac{v_n(t)}{v^*} \right)^\delta \right]$$

Equation 43.

$$\dot{a}_{break} = \left(\frac{\Delta x_{min}^*(v_n(t), \Delta v_{n,n-1}(t))}{\Delta x_{n-1,n}(t)} \right)^2$$

Equation 44.

The Gipps Model

Gipps⁽²⁾ developed the first general acceleration model that applies to both car-following and free-flow conditions. The model determines the maximum acceleration based on: driver's desired speed and minimum safe headway. The Gipps car-following model is implemented in AIMSUN. In the Gipps car-following model, the acceleration is calculated for five different situations:

- Vehicle moves in its desired speed
- Vehicle in free move with speed less than the desired speed
- Vehicle is accelerating from stopping position
- Vehicle in a car-following state with constraint space headway
- Vehicle in a car-following state with active-non collision constraint

The following thresholds constrain vehicle acceleration in Gipps' model:

$$v_n^a(t+T) = v_c(t) + 2.5 a_n^{max} T \left(1 - \frac{v_n(t)}{v_n^{desired}} \right) \cdot \sqrt{0.025 + \frac{v_n(t)}{v_n^{desired}}}$$

Equation 45.

$$v_n^b(t+T) = d_n^{max} \cdot T \cdot \sqrt{(d_n^{max} \cdot T)^2 - d_n^{max} \cdot \left[2\{x_{n-1}(t) - s_{n-1}(t)\} - v_n(t) \cdot T - \frac{v_n(t)^2}{\hat{d}_{n-1}} \right]}$$

Equation 46.

where

a_n^{max} = maximum desired acceleration, vehicle n (m/s²),

d_n^{max} = maximum desired deceleration, vehicle n (m/s²),

d_{n-1} = estimation of maximum deceleration desired by vehicle n-1 (m/s²),

S_{n-1} = effective length of vehicle,

T = reaction time,

\hat{d}_{n-1} = leader desired deceleration, and

x_{n-1} = vehicles spacing.

The VDIFF Model

In the velocity difference model (VDIFF),⁽²⁴⁾ the acceleration function consists of a term proportional to a gap-dependent “optimal velocity” $v_{opt}(s)$ and a term that accounts for velocity differences Δv as a linear stimulus. The parameter τ is the relaxation time, which describes the adaptation to a new velocity due to changes in s and v . The sensitivity parameter λ considers the crucial influence of Δv . The properties of the VDIFF are defined by the function for the optimal velocity $v_{opt}(s)$.

The parameter v_0 defines the desired velocity under free traffic conditions. The “interaction length” l_{int} determines the transition regime for the s-shaped function going from $v_{opt}(s=0) = 0$ to $v_{opt} \rightarrow v_0$ when the distance to the leading vehicles becomes large. The “form factor” β defines the shape of the equilibrium flow-density relation. VDIFF exhibits collisions for some regimes of the parameter space.

Table 3 shows the major attributes of the car-following models discussed above. The multiphase column refers to whether the model calculates the position of the subject vehicle in multiple steps. The multiregime column refers to whether the model calculates the position of the subject vehicle with different equations depending on the regime the subject vehicle falls in.

Table 3. Major similarities and differences between car-following models.

Model	Multiphase	Multiregime	Speed Diff	Space Diff	Reaction Time	Variable Acceleration
Pipes	Yes	No	Yes	No	Yes	Yes
GM	No	No	Yes	Yes	Yes	Yes
Wiedemann	No	Yes	Yes	Yes	No	No
Fritzsche	No	Yes	Yes	Yes	No	No
CA	Yes	No	No	Yes	No	Yes
SK	Yes	No	No	Yes	No	Yes
IDM	Yes	No	Yes	Yes	Yes	Yes
Gipps	Yes	No	Yes	Yes	Yes	Yes
VDIFF	No	No	Yes	No	Yes	Yes

CAR-FOLLOWING MODELS AND SIMULATION SOFTWARE

Panwai and Dia⁽²⁵⁾ compared a number of car-following models on simulation software including AIMSUN, Paramics, and VISSIM. Using data collected by an instrumented vehicle that records differences in speed and headway (leading speed, relative distance, relative speed, follower acceleration), the data was simulated using MITSIM, the Wied/Pel model, the Wied/Vis model, the Nagel/Schreckenberg model, the optimal velocity model, and the T³ model. The follower vehicle was programmed for modeling, and speed, time, and distance headways of both leader and follower were captured and compared to field measurements. The leading vehicle was simulated using the GETRAM module in AIMSUN, an external file in VISSIM, and an API interface in Paramics. Performance measures and error indicators were used to assess the fitness. Results showed similarities for psychophysical models in VISSIM and Paramics. However, the

root mean squared error and qualitative drift and goal-seeking analyses showed a substantially different car-following behavior for Paramics.

Siuhi and Kaseko,⁽²⁶⁾ Osaki,⁽²⁷⁾ Subranmanain,⁽²⁸⁾ and Ahmed⁽²⁹⁾ modified the GM model by separating the acceleration and deceleration responses. Table 4 compares the GM calibration parameters for acceleration and deceleration models.

Table 4. Comparison of parameter values in GM-based model for acceleration and deceleration models.⁽²⁶⁾

Parameter	Acceleration/ Deceleration	Ozaki (1993)	Subra- main (1996)	Ahmed (1999)	Toledo (2003)	Siuhi (2010)
Driver sensitivity	acceleration	1.1	9.21	0.0225	0.0355	1.851
Driver sensitivity	deceleration	1.1	15.24	0.0418	0.860	3.247
Speed	acceleration	-0.2	-1.667	0.722	0.291	-0.961
Speed	deceleration	0.9	1.086	N/A	N/A	1.298
Relative speed	acceleration	1	1	0.600	0.520	0.667
Relative speed	deceleration	1	1	0.682	0.143	1.243
Separation	acceleration	-0.2	0.884	-0.242	-0.16	0.667
Separation	deceleration	-0.9	-1.659	-0.151	-0.565	-1.544
Driver response time lag (s)	acceleration	N/A	1.97	N/A	N/A	0.80
Driver response time lag (s)	deceleration	N/A	2.29	N/A	N/A	0.70
Stimulus response threshold (mi/h)	acceleration	N/A	N/A	N/A	N/A	1.30
Stimulus response threshold (mi/h)	deceleration	N/A	N/A	N/A	N/A	-1.00

Siuhi and Kaseko⁽²⁶⁾ discussed findings in terms of GM calibration parameters. The study concluded that there was a need for separate deceleration and acceleration models because the responses are impacted differently by the vehicle's speed, speed difference, and separation between the vehicles.

MODELING DRIVER BEHAVIOR

Several models aim to capture driver behavior. However, little is known about the differences in car-following behavior between driver-vehicle combinations.⁽³⁰⁾ For example, calibrating the Wiedemann model, which captures driver indifference to small changes in the stimuli and allows different execution modes including emergency braking, requires estimating 18 parameters found in 17 different equations.⁽¹⁷⁾ Microsimulation software packages use a variety of car-following models, including Gipps' (AISUM2, SISTM, and DRACULA), Wiedemann's (VISSIM), Pipe's (CORSIM), and Fritzsche's (Paramics). Automated calibration parameters, such as genetic algorithms, have been used to calibrate the distribution of car-following sensitivity parameters.⁽³¹⁾

In addition, direct correlation with real driving variables is rare, and parameterization of objective behavior is still in its infancy. Ossen et al.⁽⁴⁾ studied the car-following behavior of individual drivers using vehicle trajectory data extracted from high-resolution digital images collected at a high frequency from a helicopter. The analysis was performed by estimating the parameters of different specifications of the GHR car-following rule for individual drivers. The results showed that measurement error has a larger influence than the component of noise. Their study pointed out that optimal parameters differ, and also that the suitability of a car-following model appears to be based on the individual driver data.

Brackstone et al.⁽³²⁾ used data from a series of instrumented vehicles driven by two groups. The independent variables for this study were age, approximate mileage driven per year, passivity/aggressiveness (P/A, 1–50), driver externality and internality ratings (DE and DI), and Sensation Seeking Scale V (SSSV). Among the results, the study showed that following behavior may be split in two phases (between 30 and 50 miles), and there seems to be an inverse relationship between following distance and sensation seeking.

MODEL CALIBRATION WITH LONGITUDINAL TRAJECTORY DATA

Using two models of similar complexity (number of parameters), the Intelligent Driver Model (IDM) and the Velocity Difference Model (VDIFF), Kesting and Treiber⁽⁵⁾ researched car-following behaviors of individual drivers using publicly available trajectory data for a straight, one-lane road in Stuttgart, Germany. They used a genetic algorithm to minimize the deviations between the observed driving dynamics and the simulated trajectory. One of the major findings of the study was that a significant part of the deviations between measured and simulated trajectories can be attributed to the interdriver variability and the intradriver variability (human drivers do not drive constantly over time, and their behavioral driving parameters change). The latter accounts for a large part of the deviations between simulations and empirical observations. The results showed that the calibrated parameter values of the VDIFF strongly depend on the optimization criterion, while the IDM is more robust. Also, a deterministic car-following model allows for only an average description of human driving behavior. The authors suggest using a set of time-dependent model parameters reflecting the driver adaptation process.

Ossen and Hoogendoorn⁽³⁰⁾ analyzed trajectory data and found considerable differences between car-following behaviors of individual drivers. These differences can be expressed as different optimal parameter values for reaction time and sensitivity. Special software obtained the data from digital images captured at high frequency from a helicopter in the Netherlands. Three models from the GHR model family (Chandler, Gazis, and Edie) were analyzed, with the sensitivity parameters c and reaction times parameter T derived from accelerations, relative speeds, and distances. When individual drivers were analyzed, one model outperformed the others. However, after results for all drivers were combined, no model outperformed the others. Results show that in 80 percent of cases, a relationship between the relative speed, distance, speed of the following car, and acceleration of the following car could be established.

Menneni et al.⁽³³⁾ presented a calibration methodology based on integrated use of microscopic and macroscopic data. Microscopic data from the Next Generation Simulation (NGSIM) data collection effort was utilized in range definition of calibration parameters and qualitative calibration of the VISSIM Wiedemann car-following model. NGSIM data was used to produce relative distance versus relative velocity graphs that contain vehicle following distances and speed oscillation.

Brockfeld et al.⁽³⁴⁾ tested the validity of different models through data collected from a Differential Global Positioning System-equipped vehicle on a test track in Japan. Data of the leading vehicle was fed into the model to compute the headway time series of the following vehicle. Deviations between measured and simulated headways were then used in model calibration and validation. Brockfeld et al. tested and independently calibrated 10 models (CA0.1, SK_STAR, OVM, IDM, IDMM, Newell, GIPPSLIKE, Aerde, PRITZSCHE, and MitSim). The calibration results showed that no best model exists and that the differences between individual

drivers are larger than the differences between different models. The results also showed that although FRITZSCHE and MitSim have a large number of parameters, they do not provide better results in general.

Hoogendoorn and Hoogendoorn⁽³⁵⁾ proposed a generic calibration framework for joint estimation of car-following models. The method employed relies on the generic form of most models and weights each model based on its complexity. This new approach can cross-compare models of varying complexity and even use multiple trajectories when individual trajectory data is scarce. Prior information can also be used to realistically estimate parameter values.

DRAWBACKS OF PARAMETER CALIBRATION METHODS

Punzo and Tripodi⁽³⁶⁾ addressed the problem of the calibration of the Gipps model. Traffic stream models were developed, and stationary traffic data were used to calibrate the model. By using different combinations of microscopic parameters, the study proved that the Gipps model is unable to reproduce unstable traffic phenomena. Instead, the study generated a new model to represent multiclass traffic scenarios.

Rakha et al.⁽³⁷⁾ presented a methodology to calibrate the Gipps model for the steady-state condition. Their work assumes that all drivers have similar behaviors and characteristics. Researchers converted the car-following model into its associated macroscopic traffic stream model and calibrated key macroscopic parameters (free-flow speed, speed at capacity, capacity, and jam density) using loop detector data. These calibration methods, however, can only estimate the driver's actions on average and cannot predict the driver's actions in response to varying driving situations.

ALTERNATE MODELING OF THE DRIVING PROCESS

Hamdar and Mahmassani⁽³⁸⁾ challenge the traditional discrete-event approach of most car-following models. They use hazard-based duration models, which represent the driving process as a continuous story divided into multiple episodes. Driver behavior was modeled as "car following" and "free flow episodes." As a result, changing lanes depends on the traffic conditions and the duration of time that a driver is following a given leader; the probability of changing lanes increases during the first minute and decreases after that. The authors recommended further studies because episodes for the same driver are expected to have some interdependence and drivers will surely have some influence on one another.

Hamdar et al.⁽³⁹⁾ assessed a recently formulated hazard-based duration car-following model on its performance during congested periods. To calibrate, they applied a genetic algorithm to car-following models with complex structures, including a free-flow region and a congested region. The model was calibrated against microscopic trajectory data taken from the NGSIM data. The utility of this model is derived from prospect theory of decision-making under risk. The model also considers a penalty for the risk of accidents, an expectation value, and variance and correlation time of acceleration.

DRIVING BEHAVIOR HETEROGENEITY

Chiabaut et al.⁽⁴⁰⁾ studied driving behavior heterogeneity and proposed a methodology to estimate the parameters of Newell's car-following model (NCF) at a microscopic scale using I-80 NGSIM

data. Driving behavior variability among drivers is appraised through the minimal spacing and congested wave velocity distributions. This method improves the existing calibration methods due to the measurement.

Using vehicle trajectory data collected from car-following experiments conducted on a Japanese test track with RTK GPS receivers, Ranjitkar and Nakatsuji⁽⁴¹⁾ investigated the response time of drivers. The data was further analyzed to estimate time-variant response time when considering interpersonal and intrapersonal differences in car-following behavior, such as drivers' perception time threshold in different driving conditions. Significant intrapersonal variations were observed for different speed patterns, though no definite trends were observed in such variations, emphasizing that the influence of driving conditions and intrapersonal variations in the response time cannot be ignored. It was also found that the drivers' response time under accelerating conditions is significantly different from response time under decelerating conditions.

Hamdar and Mahmassani⁽⁴²⁾ continued their studies on different driver behaviors in a simulation, instead of using the usual homogeneous traffic stream, to identify individual variations versus collective traffic patterns. The model used a utility function that assessed the change in risk when changing speeds and also the reward for increasing or reducing speed. The model was calibrated using individual trajectory data from NGSIM. Simulation was run for both homogeneous and heterogeneous traffic. The results were scattered flow-density datapoints that created a triangular shape. This is representative of instability and the hysteresis triangle, which are congestion dynamics.

INTRODUCING INCIDENT PARAMETERS WITH TRAJECTORY DATA

Different attempts have been made to simulate safety-critical events with existing car-following models. Most of the car-following models represent an environment free of safety-critical events.⁽²¹⁾ The main factors of an incident-free environment depend on the model, varying from a sensitivity term λ (GHR model), the assumption that the following vehicle will decelerate and come to a rest before hitting the leader (Gipps), the desired gap setting (IDM, IDMM), or free space to the vehicle-ahead constraints (CA). Hamdar et al.⁽²¹⁾ list safety constraints imposed in several car-following models.

Xin et al.⁽⁴³⁾ proposed a model that aims to emulate “less-than perfect” everyday driving, capturing both safe and unsafe driver behavior. The datasets used include vehicle trajectories for six crashes and four near-crashes in Minnesota as well as vehicle trajectories that were collected in a test track in Japan using GPS. The longitudinal driving tasks were described as a negative feedback control process between the external world and the driver-vehicle-unit (DVU). The authors assumed that the equations for the perceptual threshold of visual expansion rate provide “a logical and coherent mechanism for variable reaction time.” The authors determined the vehicle acceleration or deceleration action based on checking whether safety conditions are true or false. If the conditions are false, the driver is in a “subconscious” driving state, without motivation for accelerating or decelerating. The situational factors do not include multilane effects and do not incorporate personal factors and environmental factors on driver's perception-decision process. While the idea is intriguing, it still suffers from the limitations of statistical techniques because it still uses the Gipps model.

Hamdar and Mahmassani⁽²¹⁾ attempted to capture congestion dynamics and model accident-prone behaviors by calibrating and modifying (relaxing) seven car-following models: GHR, Gipps,

cellular automation (CA), SK, the intelligent driver model (IDM), the intelligent driver model with memory (IDMM), and Wiedemann. The purpose of the investigation was to relax the safety term that was designed to avoid the creation of a crash-limited following distance from the above-mentioned models and restrict the models to realistic decelerations using NGSIM data.

Chatterjee and Davis⁽⁴⁴⁾ used VTTI naturalistic data to better define the critical elements that lead to a rear-end crash. The results demonstrate the feasibility of using vehicle trajectory data to understand the interaction between individual vehicles. Assuming that the driver's behavior can be modeled as a piecewise constant series of accelerations that are input in the dynamic trajectory model, the authors try to overcome the simple brake-to-stop model. Trajectories were duplicated using the initial values for the state variables and the time history of the accelerations. Major parameters were identified: the piecewise constant acceleration, the points where the driver changes between different accelerations, the reaction time of the driver, and the critical headway. To estimate the parameters, they used a Bayesian analysis (Markov Chain Monte Carlo), and the numerical equations were solved using WINBUGS.

NATURALISTIC DATA

Naturalistic data collection is the collection of driver behavior and performance data in a natural environment. Naturalistic data collection allowed researchers to examine what happens before and during crash, near-crash, and safety-critical events for which researchers would otherwise depend on eye witness accounts and police reports. Handled properly, data regarding vehicle speed, acceleration, range, range rate, headway, time to collision (TTC), brake pedal input, and qualitative data such as pre-incident maneuvers can be used to describe driver behavior. Qualitative data such as roadway type, number of lanes, traffic density, time of day, and weather can be used to describe the driving environment. Tying these data together allows researchers to understand the conditions that exist during an event and attain baseline data. As opposed to traditional empirical approaches, naturalistic data collection process involves drivers who operate vehicles that have been equipped with specialized sensors and recording equipment (Figure 10). The drivers operate and interact with these vehicles during their normal driving routines while the data collection equipment continuously records numerous items of interest during the entire driving epoch (a prescribed driving period).

Naturalistic Data: 100-Car Database

The 100-Car Naturalistic Driving Study,⁽⁴⁵⁾ conducted through VTTI, was a light-vehicle naturalistic study in which 100 light vehicles were instrumented. The 100-Car study was unique in that it was the first study in which vehicles were instrumented with the primary goal of collecting large-scale, naturalistic driving data. A large percentage (78%) of instrumented vehicles belonged to the participants, and were driven without special instructions. Approximately 2 million vehicle miles (43,000 hours of driving data) were included in the dataset, providing a rich dataset with extreme cases of driving behavior and performance.

An in-depth analysis focusing on driver inattention conducted by Klauer et al.⁽⁴⁶⁾ used data collected in the 100-Car Naturalistic Driving Study to establish a relationship between driving behavior and crash or near-crash involvement. The study used baseline epochs reduced from the raw 100-Car database alongside the full 100-Car study event database. The baseline database consisted of a sample of 20,000 6-s segments during which the vehicle maintained a 5 mi/h or higher velocity. Crash and near-crash risks were calculated as ratios to normal, baseline driving

data for different sources of inattention. It was found that visually or manually engaging secondary tasks result in three times greater crash or near-crash risks than attentive driving. Researchers also concluded that glances that are 2 s or longer result in twice the risk of crashes or near-crashes compared to normal baseline driving.

Naturalistic Data: 34-Truck Database

VTTI also conducted The Drowsy Driver Warning System Field Operational Test (DDWS FOT)⁽⁴⁷⁾, in which naturalistic data was collected for 18 months from 103 commercial motor vehicle (CMV) drivers. Drivers from three different fleet companies participated in the study. This database yielded approximately 2.2 million driving miles and, once reduced, yielded 1,217 valid safety-critical events. The data were processed through software that flags potential events of interest based on trigger threshold values.

Naturalistic Data: 8-Truck Database

Similarly, VTTI performed another large-scale Naturalistic Truck Driving Study's (NTDS)⁽⁴⁸⁾ investigating crash risk by identifying safety-critical events related to driver sleep/rest cycles and crash countermeasures. Continuous driving performance data was collected for 100 CMV drivers from four different fleet companies. Data was collected during 4 months of normal driving routine for 4 weeks per driver. Approximately 14,600 hours of driving data and 735,000 driving miles were collected. More than 2,800 safety-critical events were identified, including 13 crashes, 58 near crashes, 1,595 crash-relevant conflicts, and 1,213 unintentional lane deviations.

Fitch et al.⁽⁴⁹⁾ conducted an examination of driver behavior leading to lane-change crashes and near-crashes using the database created from the 100-Car study. Several objectives were investigated in this analysis: classifying the types of lane change events (e.g., severity and direction of lane-change events) and analyzing striking vehicle drivers' behaviors, as well as the struck vehicle drivers' behavior and the surrounding vehicles behaviors. The study analyzed 135 lane-change events from the 100-Car study. The lane-changes were classified using three event scenarios. It was found that not all drivers used their turn signals when changing lanes to avoid a forward crash threat. It was also found that drivers visually attend to one location at a time. The study concluded that the use of drivers-assist systems can help drivers, especially in heavy traffic conditions.

Another study examining lane-change behavior conducted by Lee et al.⁽⁵⁰⁾ explored the nature and severity of lane changes in a naturalistic driving environment. Sixteen commuters were recruited to drive one of two instrumented research vehicles for 10 days each. Half of the participants' commutes were by interstate, and half commuted on a U.S. highway. Lane change identification was specified by review of the video, which sought the initiation (first lateral move by vehicle) and end points (vehicle settled in the destination lane) for each lane change. The lane change was then categorized by maneuver type, direction, severity, urgency, and success/magnitude. Eleven categories of maneuver type were identified, including slow lead vehicle, return, enter, and exit/prep exit. During the first data pass-through, all lane changes were identified, graded, and classified, totaling 8,667 lane changes categorized into one of the 11 maneuver types. The authors reported that the analysis of the full dataset resulted in many interesting findings regarding frequency, duration, urgency, and severity of lane changes in regard to maneuver type, direction, and other classification variables. The authors further studied

a subset of the full dataset, consisting of 500 lane changes using sensor data collected from the instrumented vehicles. Investigation of this additional sample included the following variables:

- Steering.
- Lateral acceleration.
- Velocity.
- Braking.
- Turn signal use.
- Eye glance location probability.
- Eye glance link value probability.
- Mean single glance time.
- Distance to forward and rearward POVs.
- TTC to forward and rearward POVs.

The sampled lane changes were more severe and were cases where a lane change collision avoidance system (CAS) was likely to help the most. The data from this study was used to provide recommendations on display location and activation criteria for designers of lane change CASs. The research provided valuable insight into the behaviors and parameters associated with lane changes, and the archived data has the potential to answer other questions related to driver behavior.

McLaughlin et al.⁽⁵¹⁾ explored the use of naturalistic driving study to investigate the potential of avoiding rear-end crashes using CASs. For this research effort, real-time data from 13 rear-end crashes and 70 rear-end near-crashes were input to evaluate three collision avoidance algorithms. The three CAS algorithms selected and modeled for evaluation were: Equations developed by Knipling et al.⁽⁵²⁾, linear regression approach described in early CAMP work⁽⁵³⁾, and an algorithm developed by Brunson et al.⁽⁵⁴⁾

The methodology provided useful guidance both in estimating benefits achieved by the algorithms and in estimating the frequency of alerts in normal driving situations. The algorithms tested were found to generate higher than acceptable alerts.

Fitch et al.⁽⁵⁵⁾ used naturalistic data collected by VTTI to study the safety benefits that may be obtained by deploying a forward collision warning system for heavy vehicles. Researchers used kinematic motion equations and Monte Carlo simulation in the study to evaluate the effects of forward collision warning (FCW) alarms on driver forward collision avoidance behavior using rear-end (RE) conflicts from a previous naturalistic driving study. The naturalistic driving dataset used was the DDWS FOT. Methods based on Volvo's Intelligent Vehicle Initiative Field Operation Test were used to identify the RE conflicts. Researchers concluded that a nationwide deployment of FCW systems in heavy vehicles could reduce the number of RE crashes by up to 21 percent. These findings were also found to be statistically significant.

McLaughlin et al.⁽⁵⁶⁾ conducted an investigation to identify factors associated with run-off-road (ROR) crashes using the 100-Car Naturalistic Driving Study database. The objective of the study was to identify conditions in which ROR crashes or near-crashes occurred and the contributing factors associated with them. ROR crashes and near-crashes from the dataset were reviewed and included in the analysis if the subject vehicle crossed a roadway boundary or if a rapid severe evasive maneuver was undertaken by the driver to avoid a crash. The researchers also conducted a random review of the entire dataset to develop a baseline database to quantify the frequency of

the studied conditions. The study found a small number of contributing factors to 122 ROR-related events chosen for the investigation. A single factor was identified as the contributing event in 75 percent of the ROR events, while multiple factors were identified as contributing factors in only 25 percent of the cases. Contributing factors of the ROR events were found to be: distraction/inattention (40 percent), secondary task distraction (36 percent) and driving-related inattention to the forward roadway (4 percent). Other factors included changes in roadway boundaries, short following distances, and lead vehicle braking.

The naturalistic data relevant to this project can be divided into two main categories: safety-critical events and car-following periods. Using naturalistic data in this project addressed several needs related to data attributes necessary to modeling and analysis as highlighted in previous research. These attributes are mainly the ability to extract individual driving episodes per driver, allowing the analysis to recognize behavior that pertains to a particular driver and determine whether that behavior changes between different episodes. The naturalistic data also allows the analysis to examine the heterogeneity of behavior among different drivers.

The car-following periods used in this research are representative of safe driving where one vehicle follows another without any conflicts arising. The safety-critical events are the cases where conflicts arise between two vehicles and evasive or braking maneuvers are taken to avoid a collision.

AGENT-BASED MODELING IN TRANSPORTATION

Agent-based modeling (ABM) is a new paradigm that describes a system from the perspective of its constituent units.⁽⁵⁷⁾ Compared to “top-down” modeling, agent-based modeling is bottom-up, with systems that are characterized by many agents or decision makers with dispersed control, many organizational levels, the ability of agents to adapt, and the use of internal models to anticipate the future.⁽⁵⁸⁾ ABM uses rules or equations to describe individual behaviors, so an agent-based model actually starts with a set of rules and uses those rules to generate data that can be analyzed.

ABM is best applied when interactions between agents are complex, nonlinear, discontinuous, or discrete. It is also useful when space is crucial and agents’ positions are not fixed, when each individual is different, and when topology of interactions is heterogeneous and complex and agents exhibit complex behavior, including learning and adaptation.⁽⁵⁷⁾

ABM has advantages based on the belief that agents are able to interact. Agent behaviors are varied, and ABM facilitates experimentation, computing, and implementing changes in the system. A circular published by the Transportation Research Board (TRB) Artificial Intelligence and Advanced Computing Application committee⁽⁵⁹⁾ provides a useful introductory overview of the concept, structure, and application of ABM in transportation, as illustrated in Figure 9.

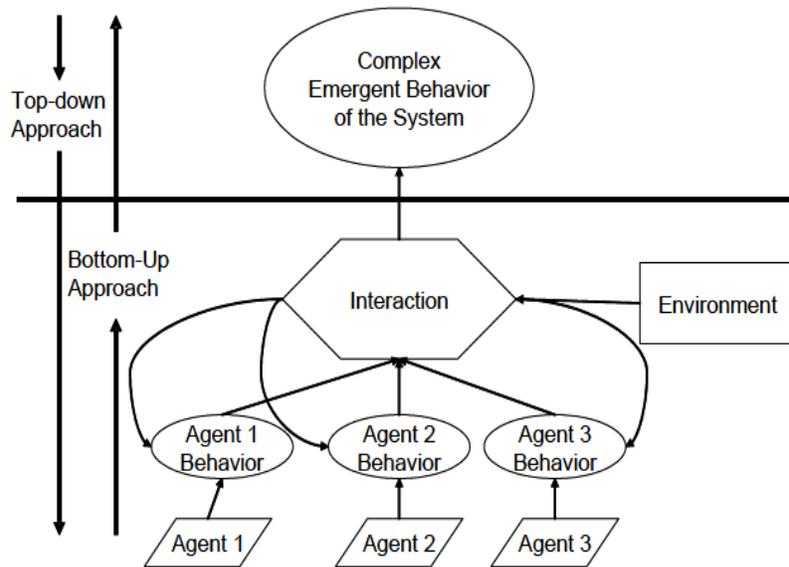


Figure 9. Illustration. “Bottom-up” modeling of system.⁽⁵⁹⁾

The ABM method is widely used in various areas of transportation, including simulation of vehicle or pedestrian flow, route choice modeling, lane-changing and car-following models, and traffic simulation. The TRB circular refers to two known simulation models as agent-based models: TRANSIMS and MATSIM.⁽⁵⁹⁾

When using ABM, types, attributes, allowable values of attributes, and initial values are defined first. Once the agents are defined, the model needs to specify the interaction rules for the agents so the agents will interact with one another and with their environment. An agent-based model can be programmed in the developer’s programming language or software package of choice.

REINFORCEMENT LEARNING METHOD

Reinforcement learning is a sub-area of machine learning in computer science concerned with how an agent is supposed to take actions in an environment so as to maximize the notion of long-term reward. The objective of reinforcement learning algorithms is to find to a policy that maps states to the actions the agent ought to take in those states. State in our case can be time from trigger, speed, car following distance, yaw angle, and environment. Reinforcement learning reinforces agents for actions performed approximately close to naturalistic actions and penalizes agents for actions that are far away from naturalistic. The only information available for learning is the system feedback, which describes the agent’s required task in terms of reward and punishment. At each time step, the agent receives a reinforcement signal according to the last action it has performed in the previous state. The problem involves optimizing not only the direct reinforcement, but also the total amount of reinforcements the agent can receive in the future. Finally, reinforcement learning should extract driving rules from a naturalistic dataset and establish similar driver-specific state action mapping rules.

The reinforcement learning method has been primarily applied to two transportation problems: finding optimal traffic control policy for intersection controllers and modeling driver behavior,

especially travel time and route choice network simulation or optimization problems. There is also a follow-up study to Abbeel et al.⁽⁶⁰⁾ in modeling helicopter control through learning from an expert pilot to simulate the trajectory and movement of helicopter control.

REINFORCEMENT LEARNING USED IN TRAFFIC ADAPTIVE CONTROL SYSTEMS

Optimization of traffic signals for efficient movement of traffic on urban streets is a challenging part of an urban traffic-control system. Adaptive system control has been used to reduce delay and congestion through an array of traffic control and management strategies (such as SCATS). The ability to exert real-time, adaptive control over a transportation process is useful for a variety of intelligent transportation systems services, including control of a system of traffic signals, paratransit vehicle dispatch control, and displaying a dynamic route guidance system.⁽⁶¹⁾

By mapping states to actions, the objective of reinforcement learning algorithms applied in adaptive control is to provide agent controllers with optimal policies, seeking the smallest queuing length and least number of stops per vehicle, according to changes in the traffic environment.

Abdulhai⁽⁶²⁾ designed a test bed and trained controllers by using Q-learning. The isolated traffic signal test-bed consisted of a simulated two-phase signal controlling the intersection of two two-lane roads. Vehicle arrivals were generated using Poisson processes with average arrival rates on each of four approaches. Q-learning agents controlled each of the four approaches with two phases of control. State information available to the agent included queue lengths and elapsed phase time. From simulation results, the performance of networked signal systems and integration with dynamic route guidance has improved. One restriction of the method is that their adaptive control method used an intersection without any turning vehicles, which is not compatible with field data.

Jacob and Abdulhai⁽⁶³⁾ used corridor control in a follow-up study to their previous work on traffic system control.⁽⁶²⁾ They proposed an automated adaptive traffic corridor control using reinforcement learning to develop a self-learning, adaptive, integrated freeway-arterial corridor control for both recurring and nonrecurring congestion. The Q-learning approach is used to provide an optimal control for a freeway corridor to determine the appropriate actions, chosen from staying on the current metering rate, increasing red time, decreasing red time, and decreasing red phase. Simulation software Paramics was used to train and evaluate an agent in an offline model within a simulated environment.

Bingham⁽⁶⁴⁾ applied an actor-critic reinforcement learning algorithm in a neuro-fuzzy traffic signal control system with the purpose of minimizing vehicular delay caused by the signal control policy. The controller received measurements of incoming traffic and chose green signal length accordingly. A neural network adjusts the fuzzy controller by fine-tuning the form and location of membership functions. Two phases were considered for any intersection. Reinforcement learning in neural networks gives credit for successful system behavior, which will be chosen more often. In simulation experiments, the simulations were run at several different traffic volumes and traffic detector locations. The result shows that different membership functions are found optimal in different traffic situations. The advantage of fuzzy control systems over traditional ones is their ability to use expert knowledge in the form of fuzzy rules and the small

number of parameters needed. However, intersections usually have four approaches and eight movements, which makes this reinforcement learning approach less persuasive.

Choy et al.⁽⁶⁵⁾ formed a multiagent architecture for real-time coordinated signal control in an urban traffic network. Three hierarchical layers of controller agents exist: intersection, zone, and regional controller. Fuzzy logic, neural networks, and regional controllers are implemented to agents. Each individual controller agent recommends an appropriate signal policy at the end of a signal phase. An online reinforcement learning module is used to update the knowledge base and rules of agents based on online data, with the objective of minimizing average delay in the zone. In all, 25 signalized intersections in a microscopic simulator were implemented in a network. Test results showed that the multiagent system improved average delay and total vehicle stop time compared with fixed-time traffic signal control. The main contribution of this paper is the formulation of a multiagent-based architecture and an online reinforcement learning module. It has shown that an agent can coordinate its local goal with zone and regional objectives autonomously.

Adam et al.⁽⁶⁶⁾ proposed a Q-learning method applied in reducing the number of vehicles traveling in dilemma zone. Algorithms used arbitrary policies to determine green termination based on the number of vehicles in the dilemma zone and whether the dilemma zone has not cleared after a period of time. A control agent develops an optimal policy by learning from the number of vehicles in the dilemma zone. The optimal policy considers the number of vehicles trapped in dilemma zone after taking action to either extend the green or not. Unlike arbitrary and stationary policies derived from existing approaches, the control policy adapts to changes in volume. Reinforcement-learning-based policy reduced the number of vehicles caught in the dilemma zone by up to 32 percent in the experimental simulation framework VISSIM.

REINFORCEMENT LEARNING IN TRANSPORTATION NETWORK PROBLEMS

Bogers et al.⁽⁶⁷⁾ focused on route choices, considering two types of learning: reinforcement-based learning and explicit belief-based learning (including memory decay). The researchers developed a model that captures learning types from parameters calibrated from a large dataset from experimental research. The model used a Markov formulation for daily updating, based on a person's belief about travel time. Reinforcement was modeled by including the latest 10 route choices in the model. Results indicate that 20 percent of perceived travel time is from the most recent experience. Furthermore, the reinforcement part of the model can make up a significant part of the route utility and therefore should be a standard component in route choice models.

Arentze and Timmermans⁽⁶⁸⁾ developed a framework for modeling dynamic choice based on a theory of reinforcement learning and adaptation because individuals develop and continuously adapt choice rules while interacting with the environment. A reward function, incremental action value functions, and action selection methods, were presented in this activity-based analysis, and memory and search played a key role in reinforcement learning. The reinforcement learning concept assumes that actions produce positive rewards if they are reinforced and have a higher probability of being repeated in future choice situations under similar conditions, and actions with negative outcomes tend to be avoided. The system assumes multistage decision making in potentially very large condition spaces and can deal with stochastic, nonstationary, and discontinuous reward functions. Arentze and Timmermans considered a hypothetical case that combined route, destination, and mode choice for an activity under time-varying conditions of the

activity schedule and road congestion probabilities. The learning and adaption was robust and most successful when the system chose a realistic reward reference level and assigned only a modest weight to new rewards in updating action values. Arentze and Timmermans' framework took exploration and adaption in choice behavior into account, but it did not take into consideration environmental effects.

Han and Timmermans⁽⁶⁹⁾ proposed a general model of interactive learning behavior that evolves toward equilibrium in strategic situations under the assumption that travelers may choose particular choice options by taking into account their expectations about behavior of other travelers. The properties of the model are examined by using numerical computer simulations. The results of the simulations support the face validity of the formulated model.

Wahba and Shalaby⁽⁷⁰⁾ proposed an operational prototype of an innovative framework for the transit assignment problem by using a learning-based approach. The proposed framework represents a passenger's learning and decision-making activities. In this prototype, a hypothetical transit network consisted of 22 routes and 194 stops developed within a microsimulation platform (Paramics). They then generated 3,000 passengers synthesized to model the transit assignment process in the peak morning periods. Reinforcement learning is used to model passengers' adaptation and account for the differences in passengers' preference in the dynamics of the transit network. Using this technique, the prototype has demonstrated that the proposed approach can simultaneously predict how passengers will choose their routes and estimate the total passenger travel cost in a congested network as well as loads on different transit routes.

REINFORCEMENT LEARNING APPLIED IN TRAFFIC KINEMATIC PROBLEM

Abbeel et al.⁽⁷¹⁾ presented the first successful autonomous completion of four aerobatic maneuvers using a real RC helicopter: forward flip, sideways roll at low speed, tail-in funnel, and nose-in funnel. First, a pilot controlled the helicopter to define a helicopter dynamics model and a reward (cost) function. Then a controller optimized for the result model and reward functions was formulated using the reinforcement learning (RL) algorithm. Differential dynamic programming (DDP) was used in this optimization process.

During the flip, the helicopter rotates 360 degrees forward around a lateral axis. In a right axial roll, the helicopter rotates 360 degrees around its longitudinal axis. In a tail-in funnel, the helicopter repeatedly flies in a circle sideways with tail pointing to center of the circle. The nose-in funnel is similar to tail-in funnel with the exception of the nose pointing to the center of the circle.

The helicopter dynamics are modeled by the apprenticeship learning algorithm method. First, the data is collected from a human pilot flying the desired maneuvers in a helicopter to learn the optimal model. Next, a controller is found to simulate a flight based on the current model test control on the helicopter. The helicopter state includes position (x, y, z), orientation, velocity ($\dot{X}, \dot{Y}, \dot{Z}$), and angular velocity ($\omega_x, \omega_y, \omega_z$). The helicopter is controlled by four separate actions: cyclic pitch (u_1, u_2), tail rotor (u_3) controls, and collective pitch angle (u_4). This paper formulates the equations for acceleration and then integrates the accelerations together to obtain helicopter states. (Superscript b indicates the body coordinate for helicopter, coefficients A, B, C, D, and E are estimated from helicopter flight data, and ω on the right-hand side are zero mean Gaussian random variables.)

The reinforcement learning method has a sextuple of components (S, A, T, H, s(0), and R), including the set of states, S; the set of actions or inputs, A; and the dynamics model, T, which is a set of probability distributions $\{P_{su}^t(s'|s,u)\}$. The probability of being in state s' at time $t+1$ given the state and action at time t are s and u . The component $s(0)$ is the initial state, and R is the reward function. A policy $\pi=(u_0, u_1, \dots, u_h)$ is the relation from the actions to the states. The expected sum of the rewards when enacting the following policy π is given by the following equation:

$$E\left[\sum_{t=0}^H R(s(t), u(t))|\pi\right]$$

Equation 47.

The optimal policy π^* is the policy that maximizes the expected sum of the rewards and is given as the following:

$$\pi^* = \operatorname{argmax}_{\pi} E\left[\sum_{t=0}^H R(s(t), u(t))|\pi\right]$$

Equation 48.

where

π^* = optimal policy,

$s(t)$ = state in time t ,

$u(t)$ = action in time t , and

$R(s(t), u(t))$ = reward function according to policy π when taking action $u(t)$ at state $s(t)$.

The RL method solves this linear quadratic regulator (LQR) control problem as the following:

$$s(t+1) = A(t) s(t) + B(t) u(t) + w(t)$$

Equation 49.

The reward for being in the state $s(t)$ and taking the action/input $u(t)$ is given by the following function:

$$s(t)^T Q(t) s(t) - u(t)^T R(t) u(t)$$

Equation 50.

where

$Q(t), R(t)$ = the positive semi-definite matrices that parameterize the reward function.

In the differential dynamic programming (DDP) design for solving LQR, the error state is defined as the following:

$$e = (x^b - (x^b)^*, y^b - (y^b)^*, z^b - (z^b)^*, x - x^*, y - y^*, z - z^*, \omega_x^b - (\omega_x^b)^*, \omega_y^b - (\omega_y^b)^*, \omega_z^b - (\omega_z^b)^*, \Delta_q)$$

Equation 51.

where

x, y, z = helicopter position

$\dot{x}, \dot{y}, \dot{z}$ = velocity, and

$\omega_x, \omega_y, \omega_z$ = angular velocity

DDP is used in the first phase of control design to find the open-loop input sequence that would be optimal in a noise-free setting. During the second phase, the DDP is applied to design the actual flight controller with a deviation from the nominal open-loop input sequence with defined inputs. The reward function contains 24 features, including the squared error state variables, the squared inputs, and the squared change in the inputs between the consecutive time steps and the squared integral of the error state variables. The reinforcement learning algorithm finds the controller using a method of apprenticeship learning through the inverse reinforcement learning algorithm. The inverse reinforcement learning system provides reward weights that are closer to the expertise of a human pilot by increasing or decreasing the weight of the features that deviate from the expert.

In the flip experiment, the cost matrices Q and R in the reward function are chosen by hand. The initial controller is oscillated in reality. The controller performs well with an increase in penalty for the changes in the inputs through consecutive time steps. The controller uses the same cost matrices for flips as it does when performing a roll. The controller outperforms the expert human pilots when it comes to both nose-in funnels and tail-in funnels. The DDP-based controller flies the helicopter successfully but only after being penalized for the controller's rapid change in actions over consecutive time steps. In conclusion, this paper uses the apprenticeship learning algorithms to determine the reward function and to learn the model. This is a two-phase controller design, first for a feasible trajectory and second for design of the actual controller.

Although this idea looks very similar to our study, the experiment is not applicable in our problem. In their approach, the agent controller is under online training by a reinforcement learning algorithm. However, in our problem, we do not have the luxury to design an agent controller to control a vehicle online in the field and collect the data. Instead, naturalistic trajectory is the only information we have access to. In fact, offline training is the only option.

Jouffe⁽⁷²⁾ proposed a realistic approach, similar to our study, to deal with various continuous traffic state-action mapping problems. In their paper, two reinforcement learning methods, Fuzzy Actor-Critic Learning (FACL) and Fuzzy Q-Learning (FQL), were used to tune online the conclusion part of Fuzzy Inference Systems. The only information available for learning is the system feedback, which describes in terms of reward and punishment the task the fuzzy agent has to realize. At each time step, the agent receives a reinforcement signal according to the last action it has performed in the previous state. The problem involves optimizing not only the direct reinforcement, but also the total amount of reinforcements the agent can receive in the future. To illustrate the use of these two learning methods, the authors first applied them to a problem that

involves finding a fuzzy controller to drive a boat from one bank to another, across a river with a strong nonlinear current. Then, the well-known Cart-Pole Balancing and Mountain-Car problems were used to compare the proposed methods to other reinforcement learning methods and focus on important characteristic aspects of FACL and FQL. The generality of the methods allowed the agent to learn every kind of reinforcement learning problem (continuous states, discrete/continuous actions, and various type of reinforcement functions). The experimental studies also show the superiority of these methods with respect to the existing methods in the literature.

CHAPTER 3. NATURALISTIC DATA PROCESSING

Naturalistic data collection is the collection of driver behavior and performance data in a natural environment. Naturalistic data collection has opened the door for tremendous possibilities in transportation research. It has allowed researchers to examine what happens in the final seconds before crash, near-crash, and safety-critical events for which researchers would otherwise depend on eye witness accounts and police reports. Handled properly, data regarding vehicle speed, acceleration, range, range rate, headway, time to collision (TTC), brake pedal input, and qualitative data such as pre-incident maneuvers can be used to describe driver behavior. Qualitative data such as roadway type, number of lanes, traffic density, time of day, and weather can be used to describe the driving environment. Tying these data together allows researchers to understand the conditions that exist during an event and attain baseline data. Naturalistic driving data collection is a powerful approach used by researchers to understand crash and near-crash causation. As opposed to traditional epidemiological and experimental/empirical approaches, this *in situ* process uses drivers who operate vehicles that have been equipped with specialized sensors as well as processing and recording equipment (Figure 10). In effect, the vehicle becomes the data collection device. The drivers operate and interact with these vehicles during their normal driving routines while the data collection equipment continuously records numerous items of interest during the entire driving epoch (a prescribed driving period).

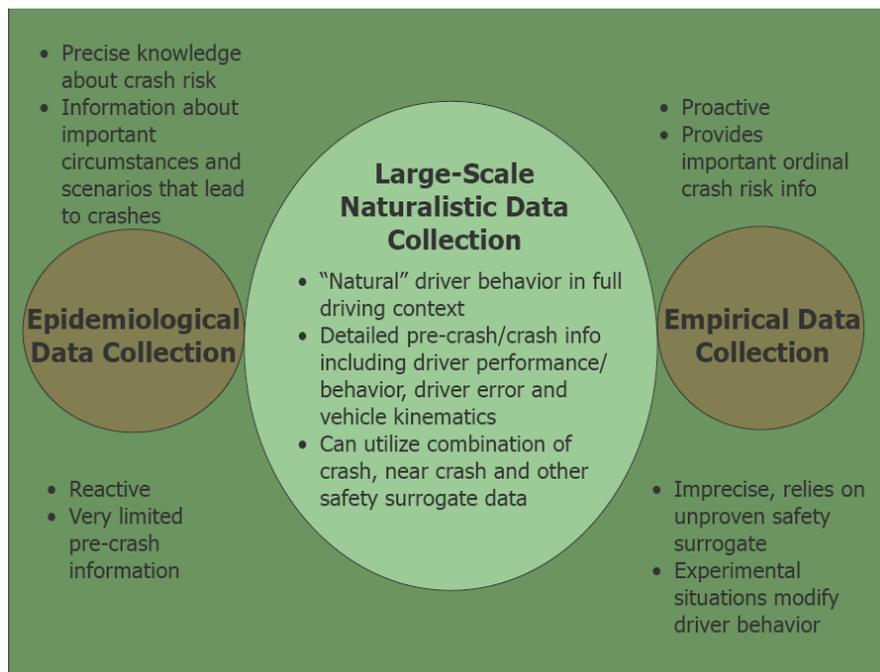


Figure 10. Diagram. Advantages of the naturalistic data collection methodology data collection and reduction.⁽⁴⁵⁾

Naturalistic data collection methods require a sophisticated network of sensors, processing, and recording systems (Figure 11). The Data Acquisition System (DAS) provides a diverse collection of both on-road driving and driver (participant, nondriving) data, including measures such as

driver input and performance (e.g., lane position, headway, etc.), four camera video views, and driver activity data. This information may be supplemented by subjective data, such as questionnaire data. A tremendous amount of data is acquired when carrying out these naturalistic studies.

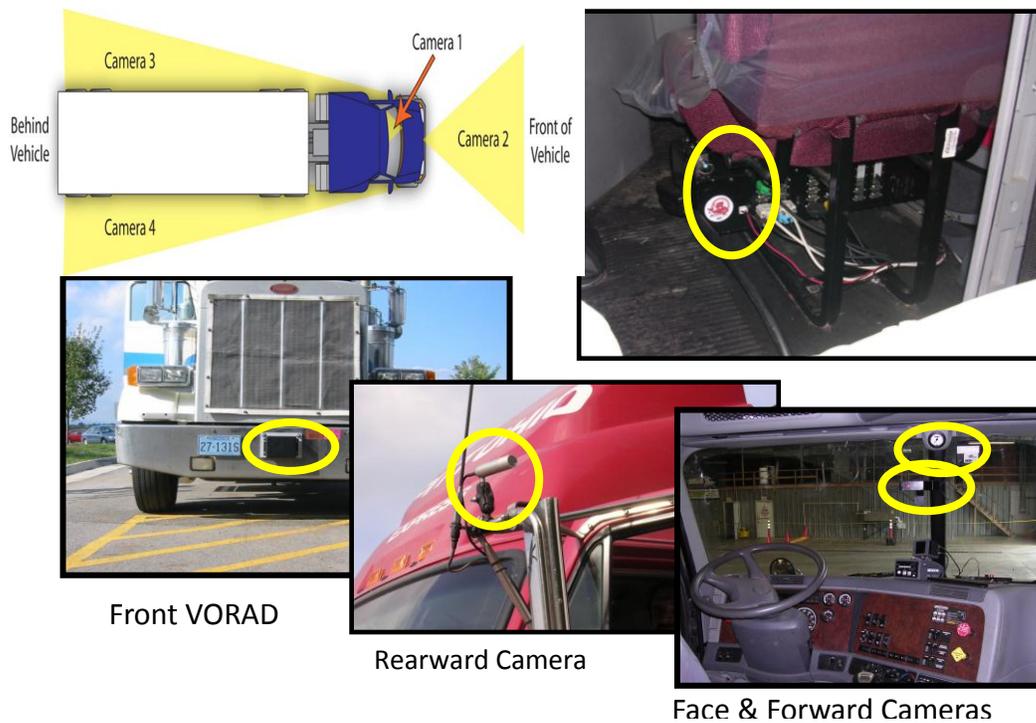
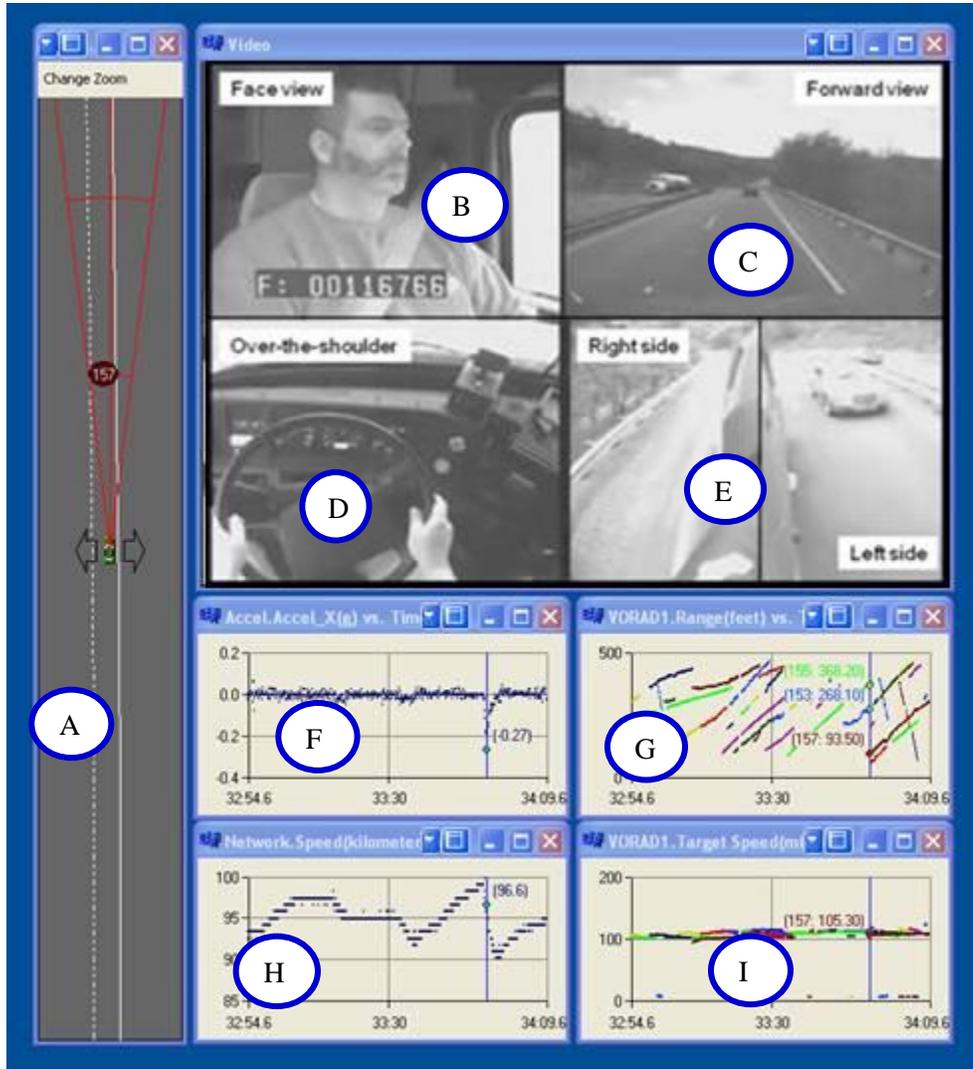


Figure 11. Illustration. Naturalistic data sensors.⁽⁴⁸⁾

Due to the tremendous amount of information collected during these procedures, naturalistic data collection methods require significant post-collection processing. VTTI processes this data using secure computing systems on isolated networks, which allows for both data processing and analysis. A storage area network is employed to store the data and to play multiple driving and video files simultaneously at several PC workstations.

Typically, the first step in the data reduction process is to identify events of interest, including crashes, near-crashes, and crash-relevant conflicts. To find events of interest, VTTI has developed a software program, Data Analysis and Reduction Tool (DART), that scans the dataset using user-defined threshold algorithms to identify notable actions (e.g., hard braking, quick steering maneuvers, short times-to-collision [TTC], and lane deviations, including median encroachments). All detected events are analyzed based on “instant replays” of video data and associated dynamic data recordings of the events. This analysis captures both the observable causal sequences leading to events as well as the conditions and correlates of event occurrence. Figure 12 presents a screen shot from the DART software interface.



A = Plan view of subject vehicle and road position based on Road Scout lane tracker.

B = Video feed recording face shot of driver, with frame number, used to synchronize video and parametric data.

C = Video feed recording forward view.

D = Video feed recording of over-the-shoulder view of driver, which is used for determining driver distractions.

E = Video feed recording both rear-left and rear-right images.

F = Data feed from X/Y Acceleration sensor indicating longitudinal acceleration as measured in gs (the range to primary target as determined by VORAD).

G = Data feed from front VORAD sensor, indicating range to primary target as determined by VORAD, versus time.

H = Network speed is the vehicle's speed.

I = Data feed showing VORAD primary target closing rate (to subject vehicle).

Figure 12. Screen Shot. Screen shot of DART software interface.

Table 5 identifies the DART components, tracking/recording mechanisms, and resulting data products. Using threshold values (triggers) developed by VTTI, the DART system was used to identify the different events.

Table 5. Data Analysis and Reduction Tool (DART) components. ⁽⁴⁵⁾

Sensors	Vehicle Network	Incident Box	Video Cameras	On-Board Components	Data Products
<ul style="list-style-type: none"> • GPS • lane tracker • x/y acceleration • front VORAD • RF sensor • seat acceleration • sound level meter • in-cab temperature 	<ul style="list-style-type: none"> • speed • distance • ignition signal • brake activation • turn signals 	<ul style="list-style-type: none"> • light level • incident pushbutton • microphone 	<ul style="list-style-type: none"> • face • forward • rear-left • rear-right 	<ul style="list-style-type: none"> • data reliability check software 	<ul style="list-style-type: none"> • truck performance data file (in .dat format) • digital video file (in .mpg format) • digital audio file (in .mp3 format)

The Drowsy Driver Warning System Field Operational Test (DDWS FOT) ⁽⁴⁷⁾ conducted by VTTI was a naturalistic data collection study that collected data for 18 months from 103 commercial motor vehicle drivers (CMV). This database yielded approximately 2.2 million driving miles and, once reduced, found 1,217 valid safety-critical events. The data was processed through software that flags potential events of interest based on trigger threshold values. ⁽⁴⁸⁾ Table 6 shows the triggers used for the DDWS FOT study.

Table 6. Trigger values used in the DDWS FOT. ⁽⁴⁷⁾

Trigger Type	Trigger Values
Longitudinal acceleration (hard braking)	Deceleration greater than or equal to -0.35 g . Speed greater than or equal to 15 mi/h. Deceleration greater than or equal to -0.5 g . Speed less than or equal to 15 mi/h.
Time-to-collision	A forward time-to-collision (TTC) value of less than or equal to 1.8 s, coupled with a range of less than or equal to 150 ft, a target speed of greater than or equal to 5 mi/h, a yaw rate of less than or equal to $ 4^\circ/\text{sec} $, and an azimuth of less than or equal to $ 0.8^\circ $. A forward TTC value of less than or equal to 1.8 s, coupled with an acceleration or deceleration greater than or equal to $ 0.35\text{g} $, a forward range of less than or equal to 150 ft, a yaw rate of less than or equal to $ 4^\circ/\text{sec} $, and an azimuth of less than or equal to $ 0.8^\circ $.
Swerve	Swerve value of greater than or equal to 3 rad/s^2 . Speed greater than or equal to 15 mi/h.
Critical incident button	Activated by the driver upon pressing a button, located by the driver's visor, when an incident occurs that he/she deems critical.
Analyst identified	Event that was identified by a data analyst viewing video footage; no other trigger listed above identified the event (i.e., longitudinal acceleration, TTC, etc.).

When analyzing data, it is critical for the researcher to have a strong dataset of baseline data or epochs. Baseline epochs are described as brief time periods (e.g., 6 s) that are randomly selected from the recorded dataset. Baseline epochs are described using many of the same variables and data elements used to describe and classify crashes, near-crashes, and crash-relevant conflicts. Examples of such variables included ambient weather, roadway type, and driver behaviors. In a recent study utilizing the dataset from the 100-Car Naturalistic Driving Study⁽⁴⁵⁾, one of the first large-scale studies, McLaughlin et al.⁽⁵⁶⁾ explain the process researchers used to utilize the 100-Car dataset in order to develop a baseline database that would quantify the frequency of different conditions and behaviors present while driving. During this process, researchers used the dataset to create a description of baseline driving. Conditions such as lighting and weather were then classified. Analysts used these random epochs to estimate miles traveled in different conditions by multiplying a participant's estimated miles traveled by the classified conditions. The sampling was stratified according to the involvement of each driver in crashes, near-crashes, and incidents of all types. From this analysis, it was determined that if a driver was involved in 3 percent of total events, then 3 percent of baseline epochs would be used from that driver's data.

NATURALISTIC TRUCK DATABASE AND TRIGGERS

The benefit of a naturalistic driving study is the ability to view the driver behavior and performance data that has been collected in a naturalistic environment. In the DDWS FOT study, a diverse set of measures were collected, including driver input/performance measures, video, actigraphy, and questionnaires. Of the 103 drivers, 24 were randomly assigned to the Control group (A^9) and 79 drivers were randomly assigned to the Experimental group (A^2B^9). A refers to the baseline condition, and B refers to the treatment condition. The superscripts refer to the number of weeks each participant drove an instrumented truck. When participants were driving in the baseline condition, the DDWS did not provide alerts. However, in the treatment condition the drivers were monitored and alerts were provided.

Three types of data were collected by the DAS: video, dynamic performance, and audio. Each driver drove 60 h in a 7-day period. Approximately 48,000 driving-data hours covering 2.2 million miles traveled were collected. Typically drivers would rotate into one of the 46 instrumented trucks, and each driver drove, on average, for 12 weeks. The DAS computer included five major components: the DDWS, sensors, vehicle network, incident box, and video cameras. Four digital video cameras were used to continuously record the driver and the driving environment. The four cameras were multiplexed into a single image. The four camera views were forward, driver's face, rear-facing-left, and rear-facing-right. A software program called Loki was developed to coordinate the data collection from the different DAS components and to integrate the data into a specific DAS output file linked to the video. More than 100 variables were collected, such as PERCLOS (percent eye closure); output; and driving performance data such as lane position, speed, and longitudinal acceleration.

Similarly, VTTI performed another large-scale naturalistic data collection study investigating crash risk by identifying safety-critical events. The objective of the Naturalistic Truck Driving Study (NTDS),⁽⁴⁸⁾ as well as Blanco et al. (in press), was to use the naturalistic driving data to investigate issues in the trucking industry such as driver sleep/rest cycles and crash countermeasures. Continuous driving performance data was collected for 100 commercial motor vehicle drivers during 4 months of their normal driving routine. Approximately 14,600 hours of driving data were collected. During this period of time, more than 2,800 safety-critical events

were identified: 13 crashes, 58 near-crashes, 1,595 crash-relevant conflicts, and 1,213 unintentional lane deviations. In addition to video and performance data, each participant was asked to wear an actigraph watch (i.e., a sleep monitor) and to fill out a daily activity log. This data provided information about the drivers' daily sleep patterns and amount of sleep, as well as other measures of interest such as time since last sleep period in reference to a critical incident. One or more crash countermeasures were also identified for each critical incident. The study recruited 100 participants from 4 participating trucking fleets. Each participant was observed for approximately 4 work weeks. Once the 4-week data collection period was over, another participant would start driving the instrumented truck.

Three forms of data were collected by the NTDS DAS: video, dynamic performance, and audio. As in the DDWS FOT, the DAS system for the NTDS consisted of five components: dynamic sensors, vehicle network, incident box, and video cameras. A lane tracker was included in the DAS system and consisted of a single analog black-and-white camera, a personal computer with a frame grabber card, and an interface-to-vehicle network for obtaining ground speed. Once installed, software automatically calibrated itself to determine the camera position. The following variables were reported:

- Distance from center of truck to left and right lane markings (estimated max error <6 inches, average error < 2 inches).
- Angular offset between truck centerline and road centerline (estimated max error <1°).
- Approximate road curvature.
- Confidence in reported values for each marking found.
- Marking characteristics, such as dashed versus solid and double versus single.
- Status information, such as in-lane or solid line crossed.

An additional camera view looking over the driver's shoulder into the lap was added to the NTDS after data analysts reported that in the DDWS FOT the drivers would often reach for an unidentifiable object outside the camera view. This additional view provided information on many potentially distracting driver behaviors. The data reduction process was similar to that of the DDWS FOT. The first step was to process data using modified trigger values (Table 7) to flag potential events of interest. The lane deviation trigger in the NTDS was not included in the DDWS FOT.

Table 7. Trigger values used in the NTDS.⁽⁴⁸⁾

Trigger Type	Trigger Values
Longitudinal acceleration (hard braking)	Deceleration greater than or equal to -0.20 g. Speed greater than or equal to 1 mi/h.
Time-to-collision	A forward time-to-collision (TTC) value of less than or equal to 2 s, coupled with a range of less than or equal to 250 ft, a target speed of greater than or equal to 5 mi/h, a gyro rate of less than or equal to $ 6^\circ/\text{sec} $, and an azimuth of less than or equal to $ 0.12^\circ $.
Swerve	Swerve value of greater than or equal to 2 rad/s ² . Speed greater than or equal to 5 mi/h.
Critical incident button	Activated by the driver upon pressing a button, located by the driver's visor, when an incident occurs that he/she deems critical.
Analyst identified	Event that was identified by a data analyst viewing video footage; no other trigger listed above identified the event (i.e., longitudinal acceleration, TTC, etc.).

Together, these CMV, heavy-vehicle naturalistic studies yielded approximately 60,000 hours and 3 million miles of continuous data, which provide an extremely rich dataset. Further analysis on databases such as these will continue to answer questions on driver behavior.

NATURALISTIC CAR DATABASE

VTTI conducted another notable naturalistic data collection study (mentioned briefly above), the 100-Car Naturalistic Study,⁽⁴⁵⁾ which instrumented 100 light vehicles. The 100-Car study was unique in that it was the first study in which vehicles were instrumented with the primary goal of collecting large-scale, naturalistic driving data. Also, 78 of the 100 instrumented vehicles were those of the participants, and the participants were given no special instructions. Approximately 2 million vehicle miles were included in the dataset, which is rich with extreme cases of driving behavior and performance. Continuous driving performance data was collected for 109 primary drivers and a total of 241 drivers during 13 months of their normal driving routine. Approximately 42,300 hours of driving data were collected. During this period of time, 15 police reported crashes, 67 nonpolice-reported crashes, 761 near-crashes, and 8,295 incidents were identified. The trigger values used to flag potential events of interest in the 100-Car study are shown in Table 8.

Table 8. Trigger values used in the 100-Car Study.⁽⁴⁵⁾

Trigger Type	Trigger Values
Lateral acceleration	Lateral motion equal to or greater than 0.7 g.
Longitudinal acceleration (hard braking)	Acceleration or deceleration equal to or greater than 0.6 g. Acceleration or deceleration equal to or greater than 0.5 g coupled with a forward TTC of 4 s or less. All longitudinal decelerations between 0.4 g and 0.5 g coupled with a forward TTC of ≤ 4 s and a corresponding forward range value at the minimum TTC not greater than 100 ft.
Forward time-to-collision	Acceleration or deceleration equal to or greater than 0.5 g coupled with a forward TTC of 4 s or less. All longitudinal decelerations between 0.4 g and 0.5 g coupled with a forward TTC of ≤ 4 s and a corresponding forward range value at the minimum TTC not greater than 100 ft.
Rear time-to-collision	Any rear TTC trigger value of 2 s or less that also has a corresponding rear range distance of ≤ 50 ft AND a rear TTC trigger value in which the absolute acceleration of the following vehicle is greater than 0.3 g.
Critical incident button	Activated by the driver upon pressing a button, located by the driver's visor, when an incident occurs that he/she deems critical.
Yaw rate	Any value greater than or equal to a plus AND minus 4 degree change in heading (i.e., vehicle must return to the same general direction of travel) within a 3-s window of time.

SAFETY-CRITICAL EVENTS

The safety-critical events for all three naturalistic data databases were identified using the triggers discussed earlier. Table 9 presents an enumeration of the types of crashes and near-crashes that were identified for each database. The safety-critical events involving animals, pedestrians, or only a single vehicle are not relevant to this project because they are more representative of special case scenarios that are not recurrent. The crashes and near-crashes involving multiple vehicles are representative of conflicts between vehicles that have a recurring nature. The safety-critical events that are not a forward or rear conflict will not have radar data, so the relative location and speed of the vehicles involved will not be known.

Table 9. Enumeration of crash and near-crash data.

	8-Truck Crashes	34-Truck Crashes	100-Car Crashes	8-Truck Near-crashes	34-Truck Near-crashes	100-Car Near-Crashes
Animal	1	4	2	1	9	10
Pedestrian					3	6
1 vehicle no objects	1		23	7	9	46
1 vehicle with objects	2	23	12	8	29	15
Multiple vehicles						
Not directly in front or behind	1		5	35	46	227
Directly in front or behind		2	27	9	22	457

Ten truck drivers and ten car drivers were selected for analysis and agent implementation based on the type of crashes and near-crashes they experienced. The drivers with the most crashes and near-crashes that were forward conflicts were selected. Figure 13 shows the demographic data of the chosen truck drivers, and Figure 14 shows their demographic data.

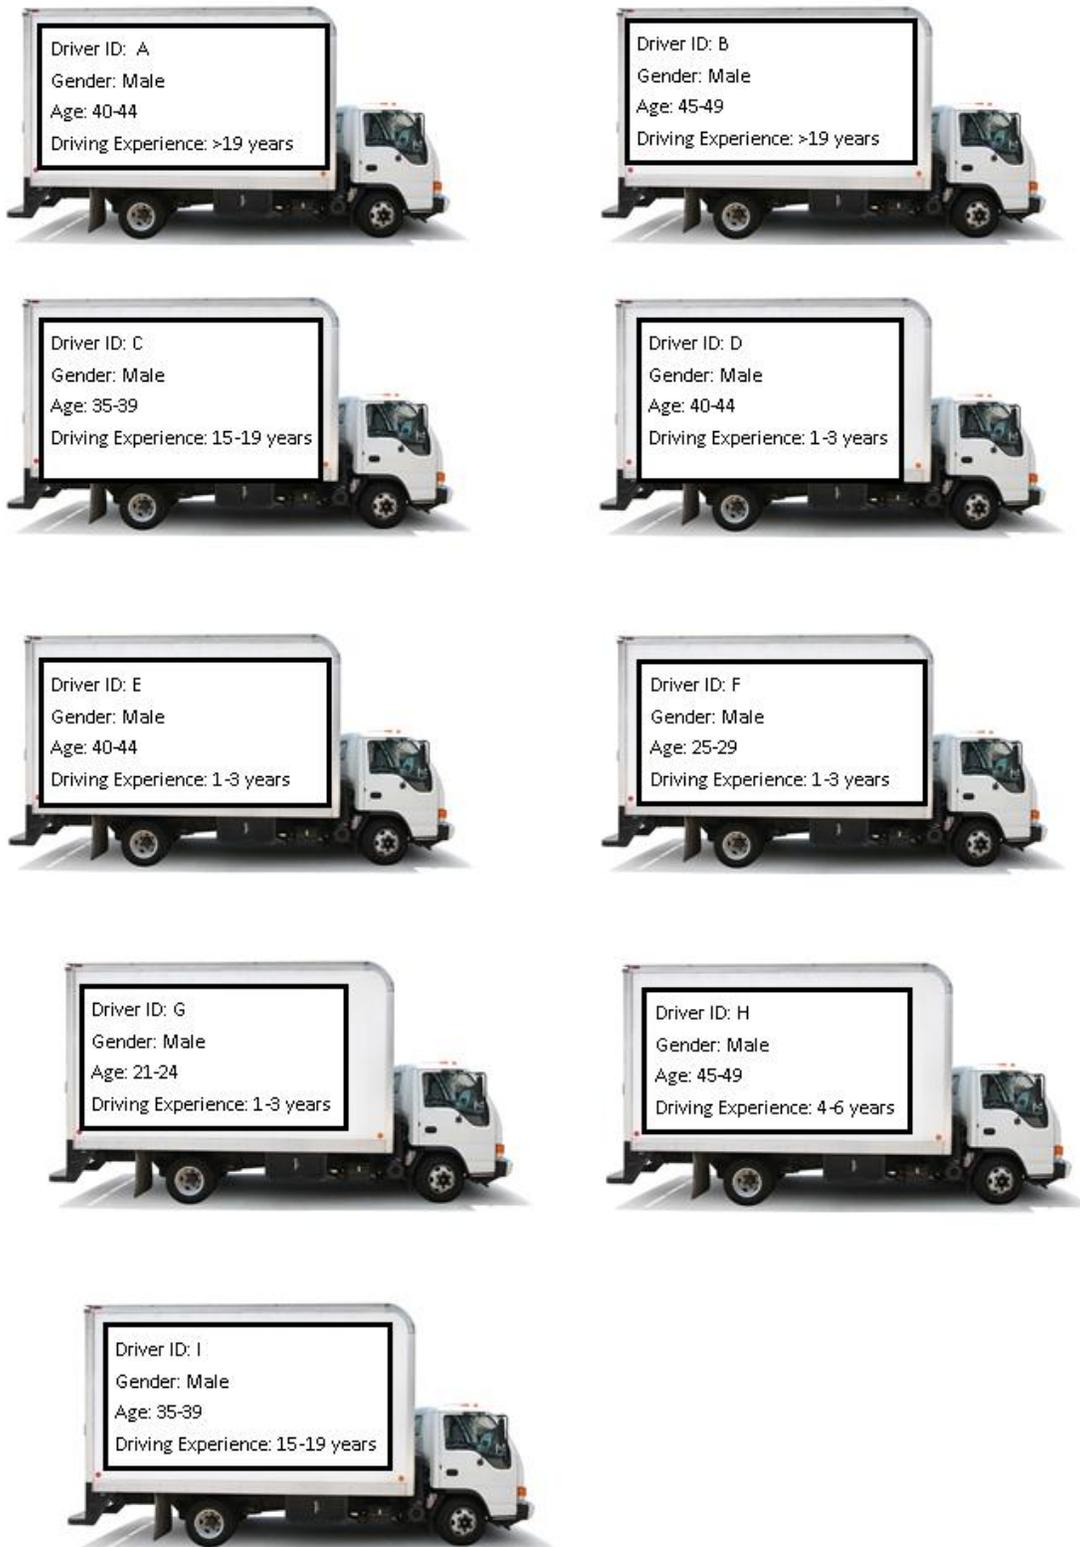


Figure 13. Illustration. Truck driver demographic data.

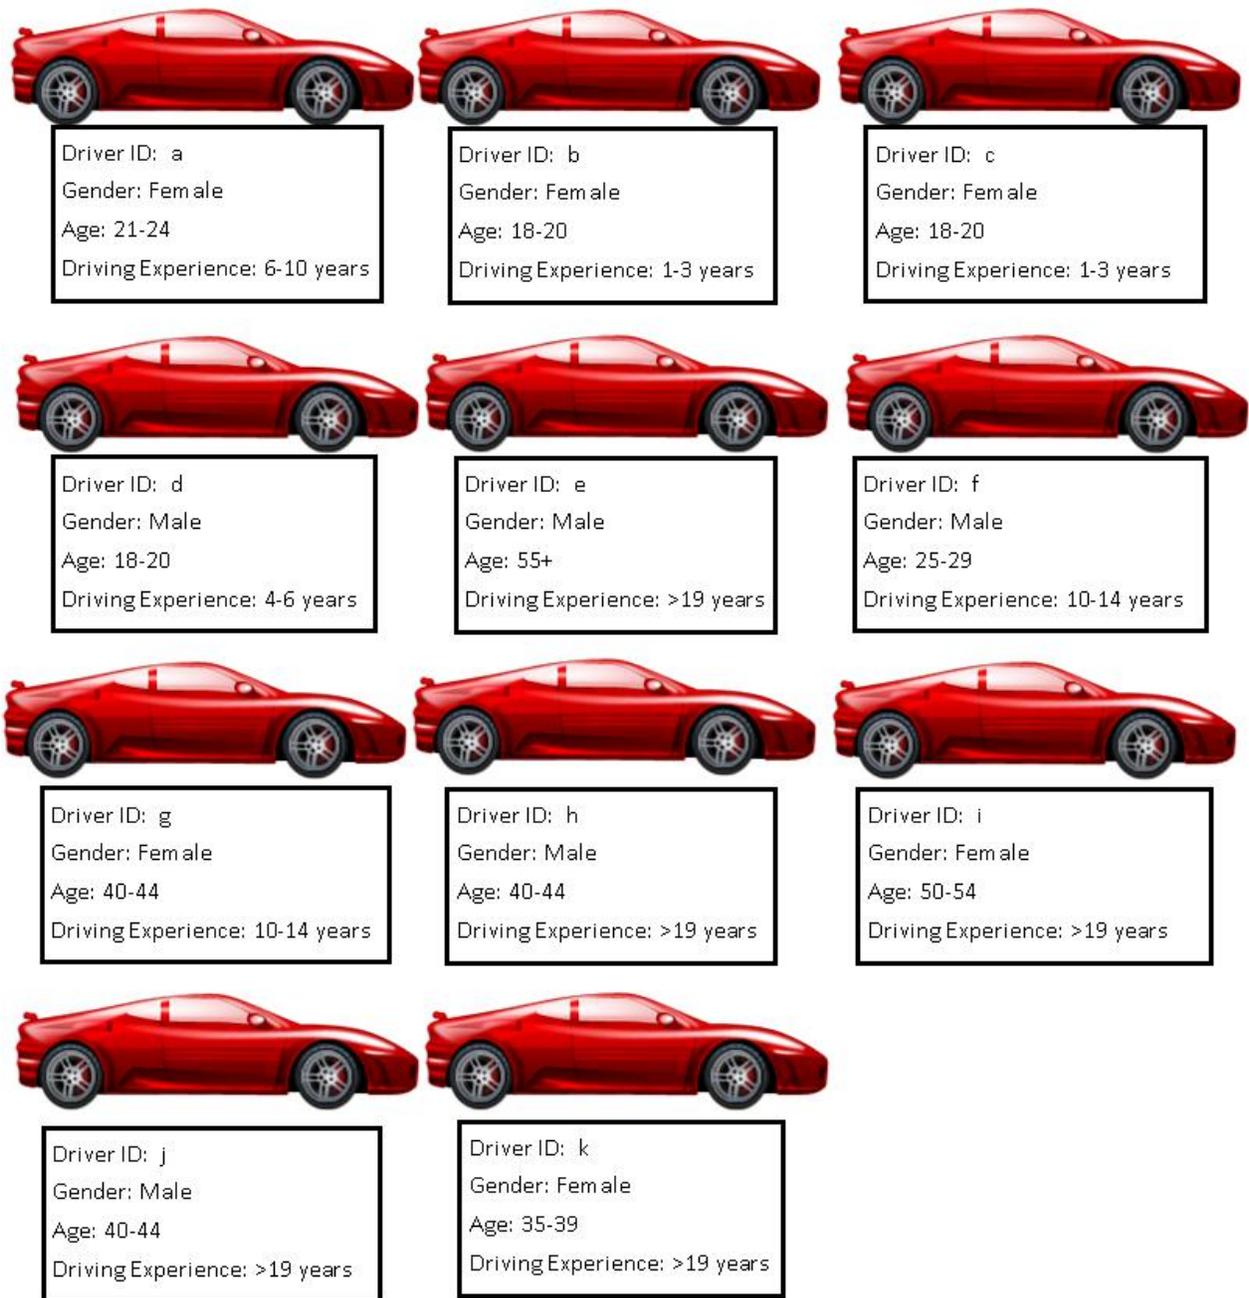


Figure 14. Illustration. Car driver demographic data.

CAR-FOLLOWING FILTERING CRITERIA

Car-following situations were automatically extracted from the enormous volume of driving data in the database to analyze car-following behavior. This process made use of the tools SQL and MATLAB to query the databases and filter them down to the needed car-following periods.

SQL is a program that is useful for efficiently accessing data stored in a database. This program allows a user to write lines of code called a query, which is then used to find all the data in the

database that matches the lines of code or filters in the query. This tool is a requirement for searching the huge volume of naturalistic data in the VTTI database. Even the data of a single driver is too large to analyze manually. For example, the data is collected at 10 Hz or 10 datapoints per second, so the 14,600 hours of data in the 8-Truck naturalistic study has 525.6 million datapoints.

An iterative process was used to filter for the car-following periods. Initial values and conditions were used, and after the events were flagged, they were reviewed in the video data to adjust the values accordingly to minimize the noise. Visual inspection of the first subsets created revealed some events that were not car-following, so additional filtering was performed to remove these events from the database.

Specifically, car-following periods were extracted automatically according to these conditions that arose as a result of the aforementioned process:

- Radar Target ID > 0.

This eliminates the points in time without a radar target detected.

- Radar Range ≤ 120 m.

This represents 4 s of headway at 70 mi/h.

- $-1.9 \text{ meters} < \text{Range} * \sin(\text{Azimuth}) < 1.9 \text{ m}$.

This restricts the data to only one lane in front of the lead vehicle.

- Speed ≥ 20 km/h.

This speed was used to minimize the effect of traffic jams but still leave the influence of congestion in the data.

- Rho-inverse $\leq 1/610 \text{ m}^{-1}$.

This limits the curvature of the roadway such that vehicles are not misidentified as being in the same lane as the subject vehicle when roadway curvature is present.

- Length of car following period ≥ 30 seconds.

All of the criteria except the length of car-following period criterion were applied in SQL. The length criterion posed difficulties in being applied in SQL. SQL evaluates each line or datapoint separately, but the criterion for length of car-following depends on the previous lines of data. MATLAB was used for this task because it offered simple and computationally efficient methods to enforce the length of car-following period criterion.

The results of these filtering criteria were verified through a video sample of 400 car-following periods, of which 396 were valid. Of the criteria listed above, it was found that the length of the car-following period criterion had the largest influence on the level of noise extracted. It was also found that increasing the length to more than 30 s reduced the number of car-following periods extracted to the point that the percentage of valid car-following periods in each sample would have little improvement over the 30-s samples. The mentioned criteria resulted in numerous car-following periods for each driver, with the exact number found being vaguely related to the amount of data available for that driver. As a rough example, a driver with around 200 hours of data would yield in excess of 1,000 car-following periods, which is more than what is needed for a single driver.

CHAPTER CONCLUSION

The naturalistic data is unique in that it captures the behavior of drivers when they are not in a test environment. This allows analyses to give results that are more related to normal, natural driving behaviors. One of the strengths of the naturalistic dataset is its sheer amount of data, which allows for more in-depth analyses. This creates a problem, however, when sorting through the data to find the data relevant to the research objectives. For this project, SQL and MATLAB were used to find car-following periods in the data, which are representative of normal behaviors during which conflicts did not arise. In contrast, the times when conflicts did arise are called safety-critical events, and these were identified in previous studies conducted by VTTI. Both datasets, safety-critical events and car-following periods, are necessary for training an agent to mimic the range of behaviors of drivers. The data from the car-following periods was used to evaluate and improve current car-following models (see Chapter 4). The event data was used to develop agents that mimic the naturalistic behavior of drivers during safety-critical events (see Chapter 5). The event data and car-following period were used in a discriminant analysis to find a way to distinguish between safe and unsafe driving behaviors (see Chapter 6). The developed agents were implemented in VISSIM (see Chapter 7). The agents implemented in VISSIM were tested and compared to the original simulation (see Chapter 8).

CHAPTER 4. CALIBRATION AND EXTENSION OF EXISTING CAR-FOLLOWING MODELS—THE WIEDEMANN MODEL

MODEL CALIBRATION

Wiedemann Car-Following Model

The Wiedemann car-following model was originally formulated in 1974 by Rainer Wiedemann.⁽¹⁵⁾ This model, also called the Wiedemann 74 model, is known for its extensive use in VISSIM, the microscopic multimodal traffic flow simulation software.⁽⁷³⁾ The simulation program chosen for the final implementation of trained agents as a result of this project was VISSIM. This makes the Wiedemann model a good illustrative example for this project. The Wiedemann model was constructed based on conceptual development and limited available data, and it must be calibrated to specific traffic-stream data.

This project used the principal ideas behind the Wiedemann model, but the exact shape or formula used in the model were updated using the naturalistic data, which is deemed to be one of the best available sources of real-world data.

Figure 15 shows the graphical form of the Wiedemann 74 model. The different thresholds are shown with a certain shape that can be amplified only during the calibration procedure. The figure shows the subject vehicle approaching a lead vehicle (ΔX decreasing due to the subject vehicle's greater speed, shown by a positive ΔV) and entering a perception area (crossing the SDV threshold) where it has to reduce speed. The subject vehicle then crosses another threshold (CLDV) where the driver reacts and reduces speed even further, entering an unconscious reaction car-following episode. The subject vehicle's driver then continues the unconscious car-following episode as long as it remains bounded by the OPDV, SDX, and SDV thresholds.

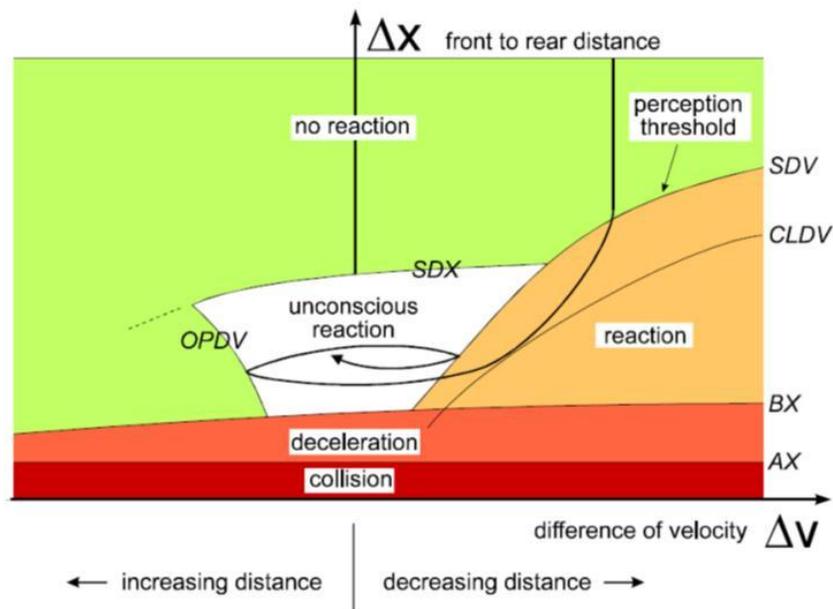


Figure 15. Diagram. Wiedemann 74 car-following logic.⁽⁷³⁾

Reconstruction of the Wiedemann Model

The reconstruction process employed has two main sections. The first section is the evaluation and adaptation of the Wiedemann model according to the naturalistic data. The second section is the addition of new thresholds that aim to represent recurrent phenomena that were found in the naturalistic data.

Calibration and Adaptation of Existing Model

The Wiedemann model uses random numbers to create heterogeneous traffic stream behavior in VISSIM. These random numbers are meant to simulate the behavior of different drivers. The naturalistic data is a perfect match for this situation because the data is collected for individual drivers. Data for four different drivers was selected and processed to reconstruct the Wiedemann car-following model.

The reconstruction began by analyzing the equations for the different thresholds and different accelerations. This coincided with the design of a calibration framework tailored specifically to the naturalistic data. The analyzed and adjusted equations were implemented into the calibration framework, a Microsoft Excel spreadsheet. Microsoft Excel 2010 was chosen for this task due to the new evolutionary algorithm included in Excel's Solver add-in. The framework expressed the logic of the Wiedemann model as a series of state transitions. The states are defined by the different thresholds, and each state has an equation or parameter for the acceleration. The optimization function was simply the minimization of the error between the values calculated in the Wiedemann model and the values obtained directly from the data.

The starting point for the Wiedemann model is the desired distance between stationary vehicles. The value calculated by Equation 52 is used in the calculations for the other thresholds. When analyzing this equation with an individual driver in mind, the calibration parameters can be combined with a driver-dependent random number to create a single calibration parameter as shown in Equation 52.

$$AX = L_{n-1} + AXadd + RND1_n \times AXmult$$

Equation 52.

where L_{n-1} is the length of the lead vehicle, $AXadd$, $AXmult$ are calibrated parameters, and $RND1_n$ is a normally distributed driver-dependent parameter.

$$AX = L_{n-1} + AXadd$$

Equation 53.

The desired minimum following distance threshold is calculated using Equation 54 and Equation 55. The calibration parameters and random number can be combined to produce Equation 56.

$$ABX = AX + BX$$

Equation 54.

$$BX = (BXadd + BXmult * RND1_n) * \sqrt{v}$$

Equation 55.

where BXadd and BXmult are calibration parameters and v is the minimum of the speed of the subject vehicle and the lead vehicle.

$$BX = BXmult * \sqrt{v}$$

Equation 56.

The maximum following distance is calculated using Equation 57 and Equation 58. The second equation includes calibration parameters and random numbers, but it mainly serves as a multiplier to BX. Since it serves as only a multiplier, the first equation can be condensed into Equation 59.

$$SDX = AX + EX * BX$$

Equation 57.

$$EX = EXadd + EXmult * (NRND - RND2_n)$$

Equation 58.

where EXadd and EXmult are calibration parameters, NRND is a normally distributed random number, and NRND2_n is a normally distributed driver-dependent parameter.

$$EX = EXmult$$

Equation 59.

The perception threshold marks the point at which a driver will begin to react to the lead vehicle. This threshold is calculated by Equation 60. Equation 53 and 61 are needed to calculate Equation 60.

$$SDV = \left(\frac{\Delta x - L_{n-1} - AX}{CX} \right)^2$$

Equation 60.

$$CX = CXconst * (CXadd + CXmult * (RND1_n + RND2_n))$$

Equation 61.

where L_{n-1} is the length of the lead vehicle; CXconst, CXadd, and CXmult are calibrated parameters; and RND1_n and RND2_n are normally distributed, driver-dependent parameters.

When looking at Equation 61 with a specific driver in mind, the equation can be reduced by combining all of the variables that would have a constant value. This collapses the equation into Equation 62, with only one input variable.

$$CX = CX$$

Equation 62.

The reaction curve marks the location of a second acceleration change point while the subject vehicle is still closing on the lead vehicle. In VISSIM, this threshold is assumed to be equivalent to the perception threshold. Due to that similarity, the equation used for the reaction threshold (Equation 63) is derived from Equation 60.

$$CLDV = \left(\frac{\Delta x - L_{n-1} - AX}{CLDVCX} \right)^2$$

Equation 63.

where CLDVCX is a calibrated parameter specific to one driver.

The OPDV (opening difference in velocity) curve is primarily a boundary to the unconscious reaction region. It represents the point where the driver notices that the distance between his or her vehicle and the lead vehicle is increasing over time. When this realization is made, the driver will accelerate to maintain desired headway. This threshold is calculated using Equation 64.

$$OPDV = CLDV * (-OPDVadd - OPDVMult * NRND)$$

Equation 64.

where OPDVadd and OPDVMult are calibrated parameters and NRND is a normally distributed random parameter.

When considering only one driver, the equation changes to Equation 65.

$$OPDV = CLDV * OPDVMult$$

Equation 65.

The Wiedemann model reuses the perception threshold as a boundary to the unconscious reaction region. This is another point at which the driver notices that the distance between his or her vehicle and the lead vehicle is decreasing over time. However, this second use of the threshold is used when the subject vehicle is already engaged in following the lead vehicle. In our reconstructed model, this reuse of the perception threshold was given its own equation to provide hysteresis control and to evaluate the adequacy of reusing the perception threshold. Equation 66 is the same as Equation 60, but with a different calibrated parameter.

$$SDV2 = \left(\frac{\Delta x - L_{n-1} - AX}{CX2} \right)^2$$

Equation 66.

where CX2 is a calibrated parameter.

The first state is the free-driving regime, in which the subject vehicle is not reacting to a lead vehicle and is traveling at a desired speed or accelerating to a desired speed. The free-driving regime is defined as the area above the perception threshold and the maximum following distance threshold. If the subject vehicle enters the free-driving regime, the subject vehicle will then accelerate until the desired speed is reached. The value for this acceleration is calculated using Equation 67 and Equation 68. The first equation relates the maximum speed to the current speed times the second equation, and it calculates an acceleration value accordingly to reach the maximum speed. Equation 69 is derived when the acceleration ends at the desired speed and not the maximum speed. Also, Equation 70 is derived from the fact that Equation 68 will reduce to a constant value.

$$b_{max} = BMAXmult * (v_{max} - v * FaktorV)$$

Equation 67.

where BMAXmult is a calibration parameter and v_{max} is the maximum speed of the vehicle.

$$FaktorV = \frac{v_{max}}{v_{des} + FAKTORVmult * (v_{max} - v_{des})}$$

Equation 68.

where FAKTORVmult is a calibration parameter and v_{des} is the desired speed.

$$b_{max} = BMAXmult * (v_{des} - v * FaktorV)$$

Equation 69.

$$FaktorV = FAKTORVmult$$

Equation 70.

The approaching regime occurs when a vehicle in the free-driving regime passes the perception threshold. This vehicle will then decelerate according to Equation 71.

$$b_n = \frac{1}{2} \frac{(\Delta v)^2}{ABX - (\Delta x - L_{n-1})} + b_{n-1}$$

Equation 71.

The closely approaching regime occurs only when a vehicle in the approaching regime passes the closing difference in velocity threshold. In VISSIM, this regime is ignored, so the deceleration is still calculated by Equation 71.

The deceleration following regime occurs as a result of a vehicle in the approaching or closely approaching regime passing the perception threshold, or a vehicle in the acceleration following regime passing the second perception threshold. When a vehicle enters the deceleration following regime, the acceleration is calculated by the negative of Equation 72. When considering one driver, this equation can be reduced to the one in Equation 73.

$$b_{null} = BNULLmult * (RND4_n + NRND)$$

Equation 72.

where BNULLmult is a calibration parameter and RND4_n is a normally distributed driver-dependent parameter.

$$b_{null} = b_{null}$$

Equation 73.

The acceleration following regime occurs when a vehicle in the deceleration following regime passes the opening difference in velocity threshold, or when a vehicle in the emergency regime passes the minimum following distance threshold. The acceleration for a vehicle in the acceleration following regime is simply the positive value of Equation 73. If a vehicle in this regime accelerates and crosses the maximum following distance threshold, then that vehicle will enter the free-driving regime. Also, the inverse is true: a vehicle will enter the acceleration following regime from the free-driving regime if the maximum following distance threshold is passed.

The emergency regime occurs any time that the space headway is below the minimum following distance threshold. Equation 74 and Equation 75 calculate the acceleration in the emergency regime. The second equation can be reduced to Equation 76 when individual drivers are considered.

$$b_n = \frac{1}{2} \frac{(\Delta v)^2}{ABX - (\Delta x - L_{n-1})} + b_{n-1} + b_{min} * \frac{ABX - (\Delta x - L_{n-1})}{BX}$$

Equation 74.

$$b_{min} = -BMINadd - BMINmult * RND3_n + BMINmult * v_n$$

Equation 75.

where BMINadd and BMINmult are calibration parameters, RND3_n is a normally distributed driver-dependent parameter, and v_n is the speed of the subject vehicle.

$$b_{min} = BMINadd + BMINmult * v_n$$

Equation 76.

New Thresholds

Hook Following Threshold

The first and most prominent recurrent phenomenon in the naturalistic data is car-following periods that begin at low space headways, with the lead vehicle traveling at a higher speed than the subject vehicle. The low space headway and higher speed indicate the completion of a pass maneuver for the lead vehicle as it merges back into the same lane as the subject vehicle. The interesting phenomenon is that the completion of the pass maneuver initiates car-following behavior. Table 10 shows that, from a sample of the car-following periods, this phenomenon represents 45 percent of the data, which is too large to ignore.

Table 10. Random sample results for number of hook following periods.

	Driver J	Driver K	Driver F	Driver L	% of Total
Regular car-following	217	131	102	256	55%
Hook car-following	166	123	142	149	45%

Figure 16 shows an example of hook car-following behavior. It is deemed hook car-following because it is initiated by the subject vehicle “hooking” onto the faster lead vehicle. The figure also shows that the vehicle moves into the unconscious reaction region and oscillates just like the regular Wiedemann model in Figure 15.

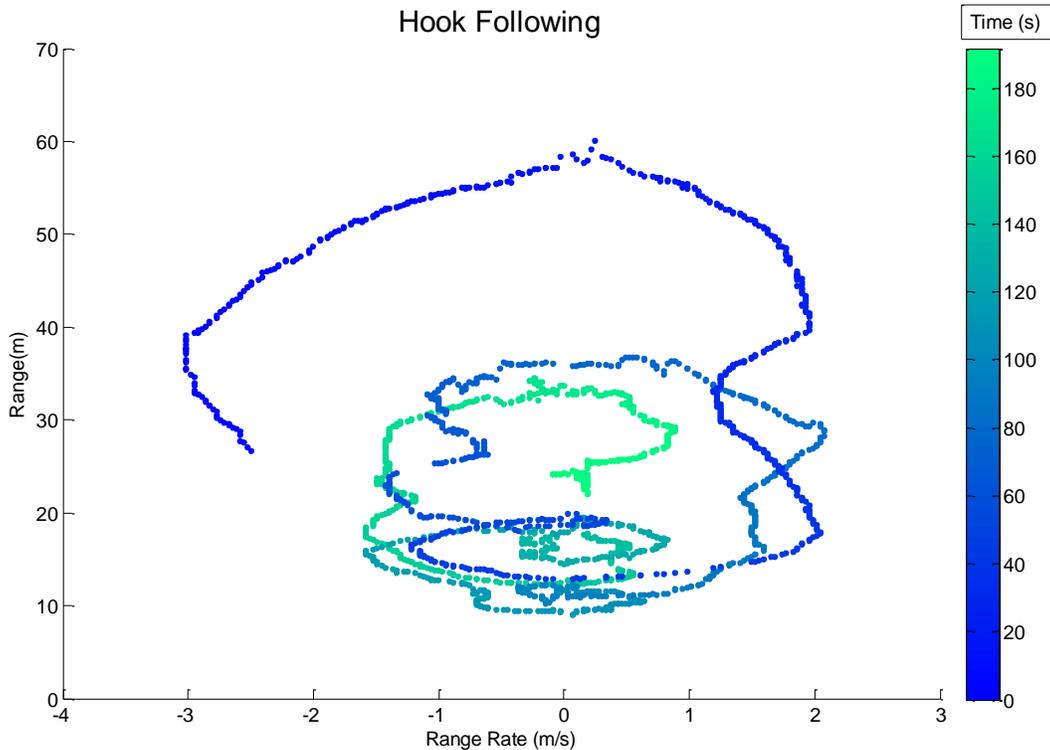


Figure 16. Graph. Example hook car-following period.

The process used was to plot two different sets of data and find the dividing line between them. The two datasets are (1) for vehicles that were not hooked onto for following and (2) for vehicles that were hooked onto and followed. The logic employed to differentiate between not hooking and hooking onto the lead vehicle operated on the fact that if a car-following period is initiated, then the subject vehicle will travel faster than the lead vehicle to catch up and follow it.

Passing Threshold

The Wiedemann model uses an external lane-changing logic for passing decisions. Trends in the naturalistic data show that this passing decision could be represented by an additional threshold to the Wiedemann model. This new threshold would operate much like the perception threshold, but the primary reaction would be to find a gap in an adjacent lane in order to pass. It will also operate like the perception threshold because the driver will need to react to the lead vehicle until an adequate gap is available.

The method employed here is the same as that employed for the hook threshold, except that in this case the subject vehicle will decelerate to follow.

RESULTS

Evaluation of the Existing Wiedemann Model

These new threshold equations were evaluated and calibrated using a genetic algorithm framework. The framework operated from state transitions, that is, where the driver would switch

states when crossing a threshold. The progression of the states was as follows: free driving, approaching, closely approaching, following decelerate, following accelerate. The emergency regime and free driving were not evaluated in order to focus on car-following behavior. The evaluation consisted of 100 car-following periods for 4 different drivers. The results of the genetic algorithm were compared to the VISSIM default values obtained from the work of Olstam and Tapani⁽³⁾ shown in Table 11. The length of the lead vehicle was evaluated as a calibration parameter to provide verification that the calibration is feasible. The results of the optimization function are shown in Table 12.

Table 11. Calibration parameters by driver.

	Default	Driver J	Driver K	Driver F	Driver L
Ln-1	4.500	4.494	5.553	4.547	5.680
AXadd	2.500	2.525	5.044	2.873	5.601
BXmult	3.000	3.041	3.405	3.421	3.435
EXmult	2.500	2.514	3.096	3.427	2.676
CX	40.000	40.066	75.042	62.405	93.131
CX2	40.000	39.031	28.978	73.258	73.026
CLDVCX	30.000	31.762	56.798	33.159	51.080
OPDVMult	-2.250	-2.512	-6.580	-1.068	-3.307
bnull	0.100	0.122	0.075	0.241	0.072
bmaxmult	0.088	0.099	0.153	0.427	0.387
FaktorVMult	0.025	0.039	0.167	0.200	0.182
bminadd	-20.000	-20.898	-13.167	-23.408	-41.721
bminmult	0.025	0.033	0.091	0.320	0.378
Vdes	40.000	42.674	84.038	88.721	16.067
FaktorV	1.000	0.940	0.522	0.507	1.959

Table 12. Root mean square error for optimization function.

	Default	Calibrated	% change
Driver J	0.784762	0.7533074	4%
Driver K	4.016006	0.9688645	76%
Driver F	3.263291	1.2555831	62%
Driver L	17.55786	1.0745801	94%

The results show that Driver J's behaviors are very similar to the VISSIM default values. The other drivers all show drastic improvement when calibrated. Also, a number of the parameters change drastically in value between drivers, which suggests that drivers should be calibrated

individually to preserve accuracy. The large variation also suggests that a generic model would be inaccurate.

New Thresholds

Hook Following Threshold

Figure 17 shows both the hooked and not-hooked reactions. As can be seen in the figure, there appears to be a dividing line between the two datasets. This same process is repeated for the other drivers, and the equations are summarized by driver in Table 13.

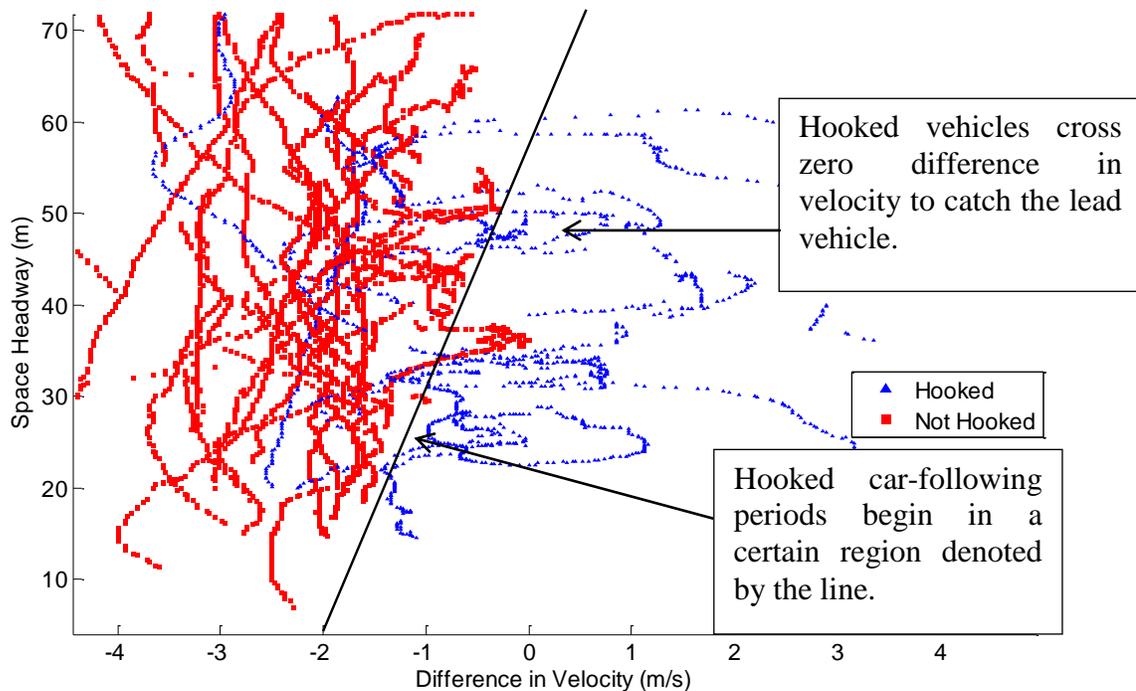


Figure 17. Graph. Hooked/not-hooked division line.

Table 13. Hook following threshold equations by driver.

Driver	Threshold Equation
J	$11.667 \cdot \Delta V + 49.667$
K	$13.822 \cdot \Delta V + 43.239$
F	$13.333 \cdot \Delta V + 46.5$
L	$15.789 \cdot \Delta V + 54.772$

Passing Threshold

The passing threshold is harder to discern than the hook threshold. Figure 18 shows data for Driver J, which is unclear until forced following is considered. In forced following, the driver intends to pass the lead vehicle, but there are no sufficient gaps allowing this maneuver. When

forced following is considered, the division between passing and not passing is clearly above any hard deceleration at low space headways. Table 14 presents the pass threshold equation for the 4 drivers.

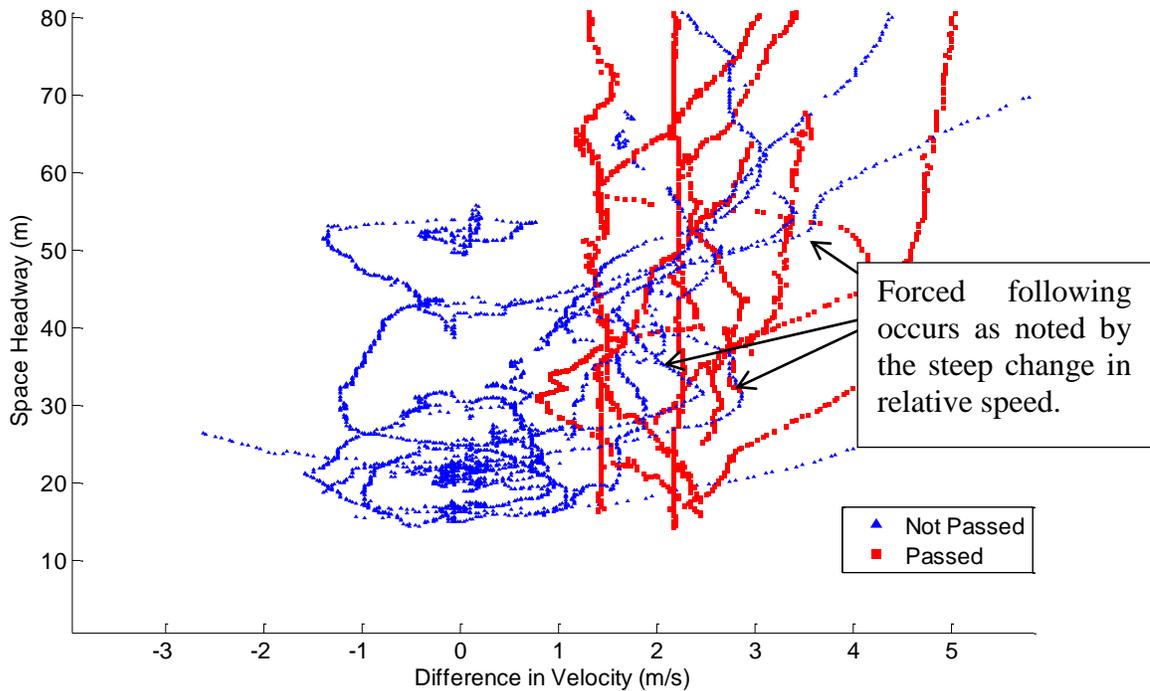


Figure 18. Graph. Pass decision curve.

Table 14. Pass threshold equations by driver.

	Pass Threshold Equation
Driver J	$0.0017\Delta x^2 - 0.0421\Delta x + 0.9981$
Driver K	$0.0007\Delta x^2 + 0.0113\Delta x + 1$
Driver F	$0.0017\Delta x^2 - 0.0672\Delta x + 1$
Driver L	$0.0002\Delta x^2 - 0.0007\Delta x + 1.5$

Reconstructed Wiedemann Model

The calibrated parameters and new thresholds were calculated for each driver and are presented in Figure 19, Figure 20, Figure 21, and Figure 22. The SDV (perception threshold) and SDV2 (second perception threshold) are very similar for three of the drivers, but very different for Driver K. For Driver K, the CLDV (closing difference in velocity) threshold is greater than the SDV2, which means that Driver K prefers to decelerate at a harder rate when approaching a vehicle within the following regime. All of the drivers have SDV2s that minimize the effect of the acceleration following regime. The ABX (minimum following distance) and SDX (maximum following distance) thresholds show similarities between Drivers K and F and between Drivers J and L.

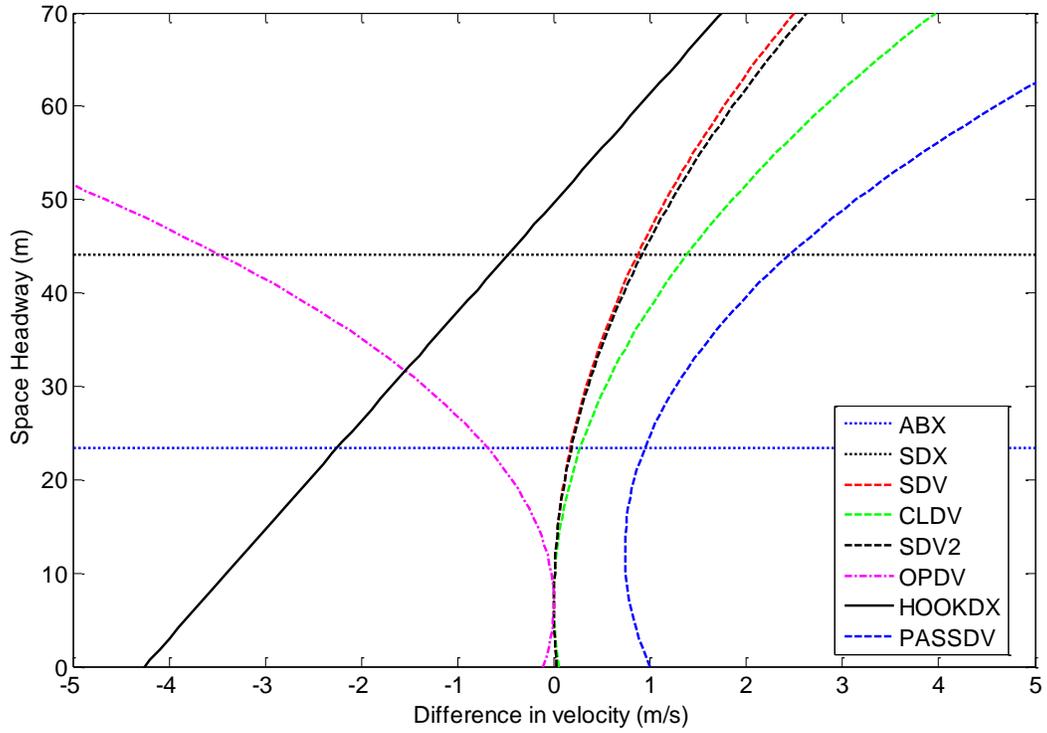


Figure 19. Graph. Thresholds for Driver J.

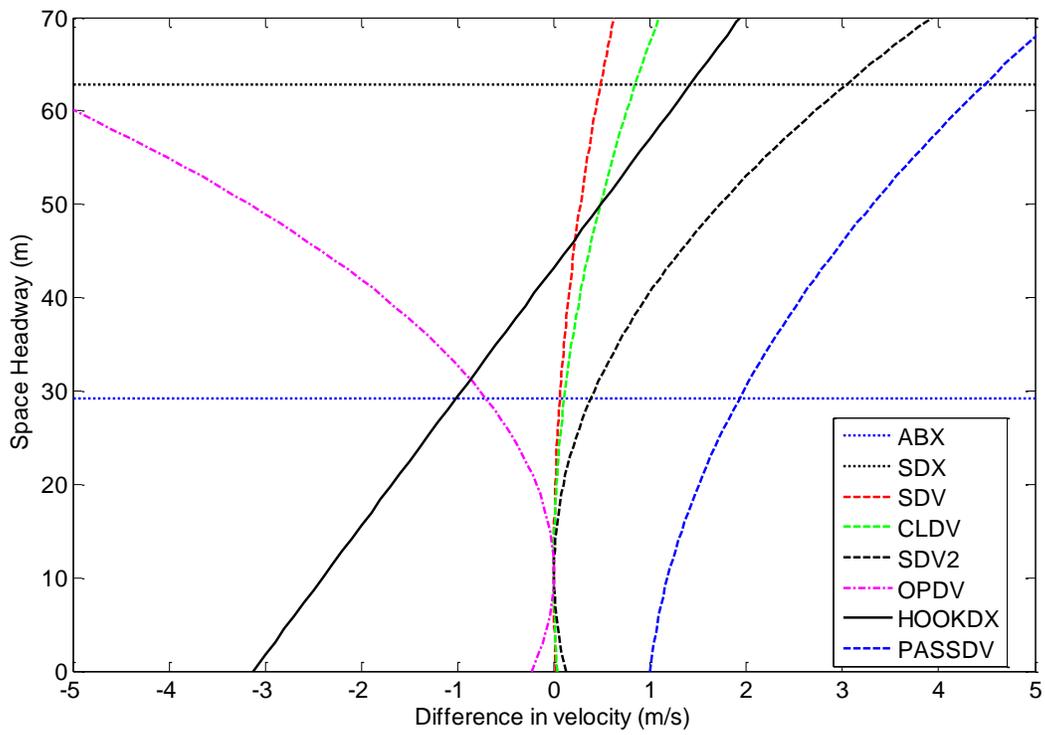


Figure 20. Graph. Thresholds for Driver K.

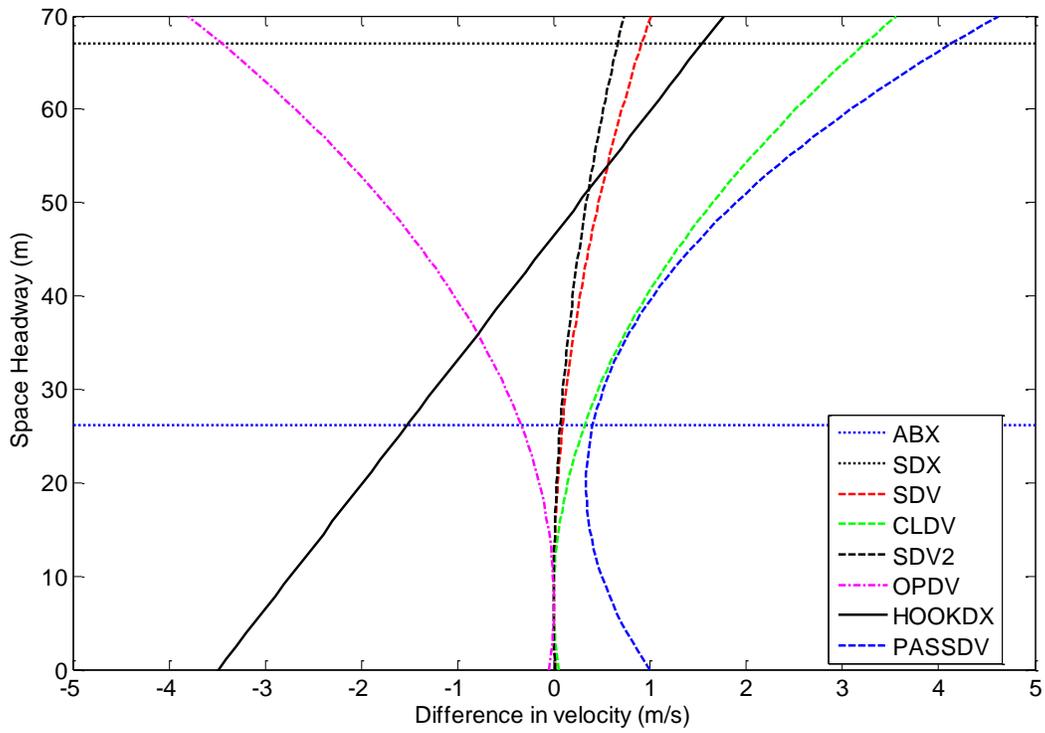


Figure 21. Graph. Thresholds for Driver F.

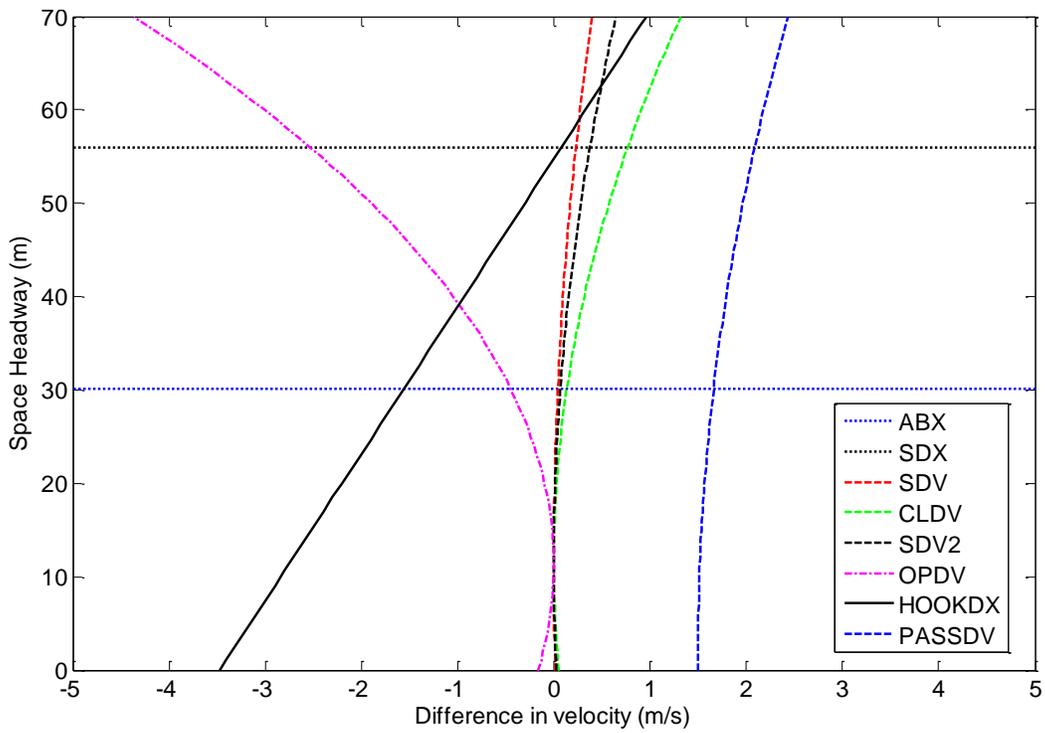


Figure 22. Graph. Thresholds for Driver L.

Conclusions

Modeling and calibrating drivers individually in this project added a new perspective to the Wiedemann model and allowed the authors to reconstruct the model by adding new thresholds. The inclusion of a hook following threshold adds value to the model by providing the ability to include a significant natural driving behavior. The addition of the passing threshold provides the model with the ability to easily transition from car-following to lane-changing behavior. This threshold also provides a way to force car-following behavior when a lane change is not possible. The addition of these new thresholds and the driver-specific equations give the Wiedemann model a better ability to represent real-world driving behavior, which is the goal of all car-following models.

HYBRID MODEL DEVELOPMENT

GHR Model

The GHR model relates the acceleration to the current speed, relative speed, and space headway as shown in Equation 77.

$$a_n(t) = cv_n^m(t) \frac{\Delta v(t-T)}{\Delta x(t-T)}$$

Equation 77.

where

$a_n(t)$ is the acceleration of the subject vehicle at time t .

$v_n(t)$ is the speed of the subject vehicle at time t .

T is the perception reaction time of the driver.

$\Delta v(t-T)$ is the relative speed at time t minus T .

$\Delta x(t-T)$ is the space headway at time t minus T .

C , l , and m are calibration parameters.

Wiedemann Model Equations

The first state in the Wiedemann model is the free-driving regime, in which the subject vehicle is not reacting to a lead vehicle and is traveling at a desired speed or accelerating to a desired speed. The free-driving regime is defined as the area above the perception threshold and the maximum following distance threshold. This regime was not included in the calibration because it does not adhere to any car-following rules.

The approaching regime occurs when a vehicle in the free-driving regime passes the perception threshold. This vehicle will then decelerate according to Equation 78. In our hybrid model, the GHR acceleration in Equation 77 replaces Equation 78. The GHR model, in this case, will have its own set of calibrated parameters specifically for approaching a lead vehicle.

$$b_n = \frac{1}{2} \frac{(\Delta v)^2}{ABX - (\Delta x - L_{n-1})} + b_{n-1}$$

Equation 78.

The closely approaching regime occurs only when a vehicle in the approaching regime passes the closing difference in velocity threshold. In VISSIM, this regime is ignored, so the deceleration is still calculated by Equation 78. The closely approaching regime would typically result in a harder deceleration than the approaching regime due to the low space headway. The GHR model accounts for this difference by using a specialized set of calibration parameters.

The deceleration following regime occurs as a result of a vehicle in the approaching or closely approaching regime passing the perception threshold, or a vehicle in the acceleration following regime passing the second perception threshold. When a vehicle enters the deceleration following regime, the acceleration is calculated by the negative of Equation 79.

$$b_{null} = BNULLmult * (RND4_n + NRND)$$

Equation 79.

where BNULLmult is a calibration parameter and RND4_n is a normally distributed driver-dependent parameter.

The acceleration following regime occurs when a vehicle in the deceleration following regime passes the opening difference in velocity threshold, or when a vehicle in the emergency regime passes the minimum following distance threshold. The acceleration for a vehicle in the acceleration following regime is simply the positive value of Equation 79. If a vehicle in this regime accelerates and crosses the maximum following distance threshold, then that vehicle will enter the free-driving regime. Also, the inverse is true: a vehicle will enter the acceleration following regime from the free-driving regime if the maximum following distance threshold is passed.

Equation 79 lacks any form of reaction to the lead vehicle. The only reaction is to switch from acceleration to deceleration once a threshold is crossed. This relies solely on the assumption that a constant small acceleration or deceleration will be sufficient to account for actions of the lead vehicle while in the following regime. The GHR model can be used to provide clarity in both of the following regimes. The two regimes were calibrated with a different set of parameters to account and test for any difference between the accelerations and decelerations of both following regimes.

The emergency regime occurs any time that the space headway is below the minimum following distance threshold. Equation 80 and Equation 81 calculate the acceleration in the emergency regime. This regime is not included in the calibration because it resembles crash and near-crash reactions more than car following and is left to the agent simulation described in Chapter 5.

$$b_n = \frac{1}{2} \frac{(\Delta v)^2}{ABX - (\Delta x - L_{n-1})} + b_{n-1} + b_{min} * \frac{ABX - (\Delta x - L_{n-1})}{BX}$$

Equation 80.

$$b_{min} = -BMINadd - BMINmult * RND3_n + BMINmult * v_n$$

Equation 81.

where $BMINadd$ and $BMINmult$ are calibration parameters, $RND3_n$ is a normally distributed driver-dependent parameter, And v_n is the speed of the subject vehicle.

Calibration Framework

The car-following data for four different drivers from the naturalistic data was used to calibrate the models by using a genetic algorithm. Ten car-following periods per driver were used in the calibration. The framework consisted of calibrating the parameters for the Wiedemann model based on the equations shown earlier. This framework was then altered such that the GHR model replaced the acceleration equations in the following regimes: approaching, closely approaching, acceleration following, and deceleration following. The GHR model was given a different set of calibration parameters (c , l , m , and T) for each of those regimes. The reaction time, T , was used as a calibration parameter to obtain a measure of the attentiveness of the different drivers.

Results

Figure 23 shows how the Wiedemann model in its current state is very reactive to the changes in the speed of the lead vehicle. This reactivity can be attributed to the fact that Equation 78 is a relative acceleration equation. Also, the recovery from a hard deceleration does not match the data in Figure 23.

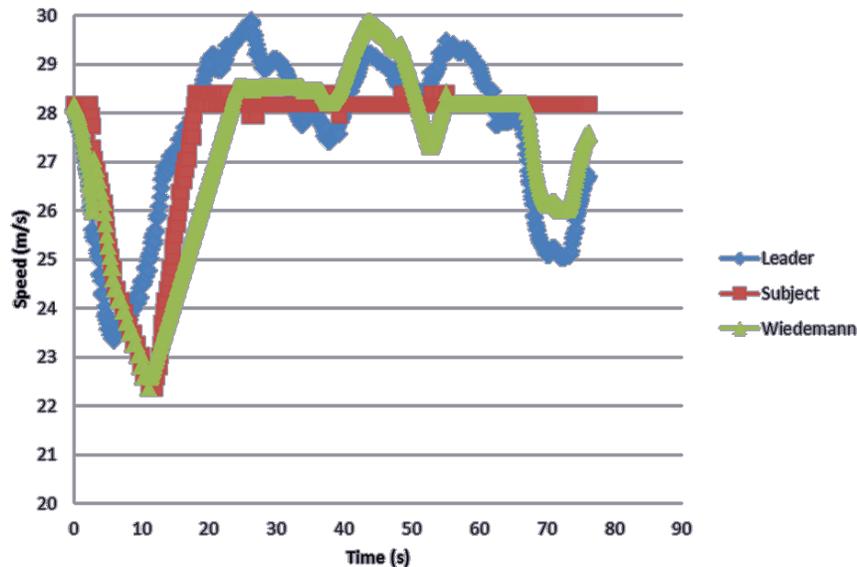


Figure 23. Graph. Driver F Wiedemann car-following period.

Figure 24 presents the same car-following episode as Figure 23, but with the GHR model integrated into the Wiedemann model. Figure 24 shows that the GHR model has removed the inaccurate reactivity to changes in the speed of the lead vehicle.

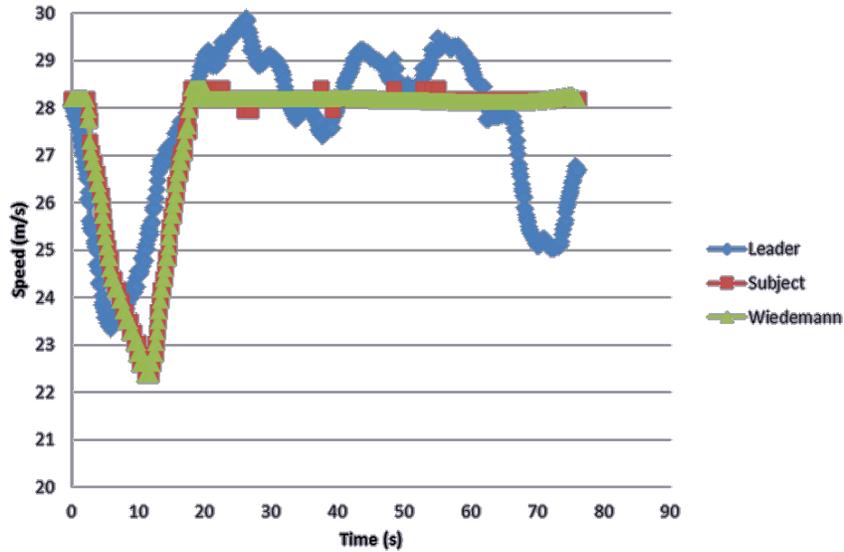


Figure 24. Graph. Driver F Wiedemann with GHR car-following period.

Figure 25 shows a comparison of the Wiedemann model and the Wiedemann model integrated with the GHR model. The GHR model corrects the recovery from a hard deceleration error seen in the Wiedemann model. The Wiedemann model appears to accelerate at a lower rate and for a longer period of time than is found in the naturalistic data.

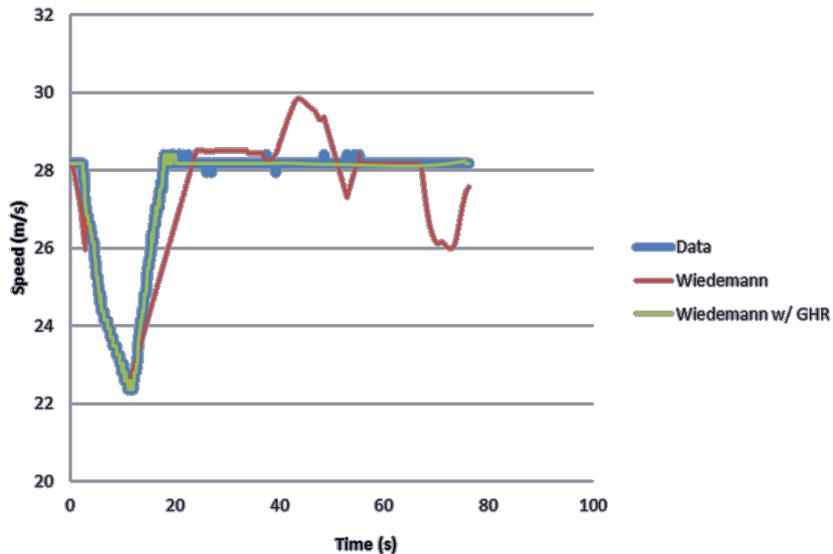


Figure 25. Graph. Comparison of Wiedemann and Wiedemann with GHR.

Figure 26 shows how the models compare in a graph showing range versus range rate. The Wiedemann model shows improper oscillation toward the end of the car-following period. Also, the Wiedemann model continues producing small oscillations at a high range when the data shows a larger oscillation moving to a smaller range at the end of the car-following period.

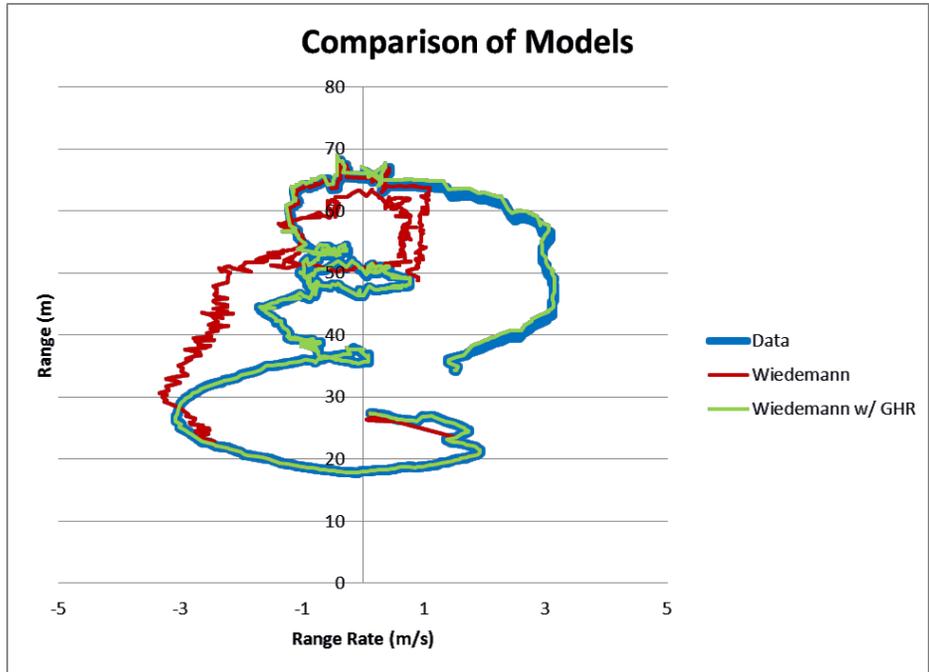


Figure 26. Graph. Comparison between models in range versus range rate.

Table 15 presents the results for the calibrated parameters of the GHR model by driver and regime. The values show variance between the drivers. The calibrated perception reaction times are low in some cases, which indicates that, in these cases, the thresholds serve as decision points and that actions from that point on are based on that decision.

Low reaction time values in the following regimes (OPDV and SDV2) indicate more active and focused following behavior. High reaction times in the same region indicate more relaxed and inattentive following behavior. This is indicated by the fact that the stimulus for the reaction is either a change in relative speed or a change in the space headway.

Table 15. GHR calibrated parameters by driver and regime.

	Driver J	Driver K	Driver F	Driver L
c SDV	0.064	0.935	0.446	0.434
m SDV	0.704	0.843	1.158	0.954
l SDV	2.135	1.997	2.046	1.933
T SDV	0.511	0.170	1.042	0.143
c SDV2	0.688	1.522	0.617	0.496
m SDV2	0.655	-0.592	0.734	0.337
l SDV2	2.613	2.565	1.821	2.002
T SDV2	0.227	0.511	0.498	0.505
c CLDV	0.130	0.532	0.602	0.706
m CLDV	0.880	-0.007	0.900	0.708

l CLDV	2.854	2.033	2.960	1.685
T CLDV	0.418	0.235	0.602	0.148
c OPDV	1.381	1.231	1.218	1.126
m OPDV	-0.521	-0.375	-0.064	-0.193
l OPDV	1.823	2.617	3.200	2.001
T OPDV	0.191	0.247	0.835	0.472

Table 16 presents the root mean square error values for both the Wiedemann model and the Wiedemann model integrated with the GHR model. The results show large improvement for two of the drivers but little improvement for the other two drivers. This indicates that the Wiedemann model calibrates to two of the drivers better than to the other two drivers. This means that the equations of the Wiedemann model reflect the behaviors of two of the drivers, but they do not accurately reflect the behaviors of the other two drivers.

Table 16. Root mean square error by driver and model.

	Wiedemann	Wiedemann with GHR	% improvement
Driver J	0.6684	0.6269	6%
Driver K	0.6042	0.4628	23%
Driver F	0.5980	0.3423	43%
Driver L	0.8106	0.7715	5%

Conclusions

The car-following data for four different drivers from the naturalistic data was used to calibrate the models using a genetic algorithm. Ten car-following periods per driver were used in the calibration. The framework consisted of calibrating the parameters for the Wiedemann model based on the equations for the thresholds and accelerations in the Wiedemann model. This framework was then altered such that the GHR model replaced the acceleration equations in the following regimes: approaching, closely approaching, acceleration following, and deceleration following. The GHR model was given a different set of calibration parameters for each of those regimes. The reaction time, T , was used as a calibration parameter to obtain a measure of the attentiveness of the different drivers. The combination of the Wiedemann model and the GHR model presents advantages when calibrating to the car-following behaviors of individual drivers. The hybrid Wiedemann-GHR model calibrated to four individual drivers' results with 5 percent to 43 percent less error than the Wiedemann model alone. Future research should look into the transition between following and free-driving regimes and the driver actions in the emergency regime.

MODEL COMPARISON

The method employed by this research effort selected, calibrated, and compared various car-following models. The models chosen for comparison are GHR, Wiedemann, Fritzsche, Gipps,

IDM, and VDIFF. These models were chosen because they represent a variety of the types of commonly used car-following models. Each model was calibrated by the use of a genetic algorithm. A genetic algorithm was used because of its ability to adequately and accurately find the optimal solution when multiple parameters are present, as in some of the models.

These models were calibrated using more than 100 car-following periods for each of the 4 truck drivers. These four truck drivers were chosen because they represent a wide spectrum of the population. When classifying by the number of driving conflicts each driver experienced, Driver 49 represents a driver with a low number of conflicts, Drivers 48 and 97 represent drivers with an average number of conflicts, and Driver 64 represents a driver with a high number of conflicts. Table 17 shows the root mean squared error that resulted from calibrating these car-following models to the four chosen drivers.

Table 17. Root mean squared error for each model by driver.

	Driver J	Driver K	Driver F	Driver L
GHR	0.9294	0.9533	1.1892	4.3815
Gipps	1.3894	1.2036	2.3112	3.8024
IDM	1.2260	1.2215	1.9211	1.0985
VDIFF	1.1082	1.0848	1.1990	1.2454
Wiedemann	0.7541	1.0916	1.4735	1.0764
Fritzsche	2.2094	1.3413	1.4332	1.1242

Table 18 shows the sum of squared error for each model by driver. The rightmost column presents the root mean squared error if the sum of squared error is summed for all the drivers and then converted to a root mean squared error. This provides a means of comparing the different models, if a single model were to be used for all of the drivers. The results show that the Wiedemann model results in the least root mean squared error, with the velocity difference model (VDIFF) not far behind.

Table 18. Sum of squared error for the models by driver with a root mean squared error of the total.

	Driver J	Driver K	Driver F	Driver L	Total	RMSE
GHR	78774	104980	177220	1752741	2113714	2.2345
Gipps	176030	167333	669436	1320099	2332898	2.3475
IDM	137063	172342	462506	110167	882078	1.4435
VDIFF	111990	135938	180169	141616	569713	1.1601
Wiedemann	51749	108426	197567	105430	463172	1.0460
Fritzsche	445138	207813	257427	115386	1025765	1.5566

Table 19 shows the parameter values of the GHR model for each driver. The GHR model showed the least root mean squared error for two of the drivers, but with the inclusion of the other drivers, this model shows a greater root mean squared error than some of the other models. This

means only that the GHR model can accurately mimic some drivers, but it is not sufficient to model all the drivers presented here. The structure of the model causes this to occur because some of the drivers exhibit behavior that is more complex than the GHR model can handle.

Table 19. GHR parameter values by driver.

	Driver J	Driver K	Driver F	Driver L
c	-3.553	-0.959	-0.542	0.300
m	-0.253	0.264	1.068	0.000
l	0.644	0.576	1.000	1.000
T	1.791	0.272	1.087	1.000
RMSE	0.929	0.953	1.189	4.381

Table 20 presents the calibrated parameters of the Gipps model for each of the drivers. The Gipps model does not calibrate as well as the other models, as shown by the root mean squared error. The Gipps model shows a large increase in error for Drivers F and L as compared to Drivers J and K. Table 20 suggests that the Gipps model is not as adequate as the other models when representing multiple drivers. The Gipps model is originally a two-regime macroscopic traffic flow model. That characteristic is obvious when the Gipps model is compared to other microscopic models, as this model tends to represent steady-state behavior.

Table 20. Gipps parameter values by driver.

	Driver J	Driver K	Driver F	Driver L
b	0.2105	0.2740	-0.0164	-0.4231
b'	6.3986	9.7948	0.5938	0.0916
Ln-1	90.1440	19.3718	10.5249	3.0981
Un	92.1683	135.6518	26.4435	25.0118
an	7.3557	7.9477	-0.0004	0.5487
T	1.9306	0.9083	0.9974	0.2282
RMSE	1.3894	1.2036	2.3112	3.8024

Table 21 presents the calibrated parameters of the intelligent driver model for the four drivers. The intelligent driver model shows a root mean squared error very similar to that of the Wiedemann model for Driver L, but greater root mean squared error values for the other drivers. In Table 21, the intelligent driver model shows a root mean squared error that suggests it is better at representing multiple drivers than some other models, but it is not as adequate as they are. This means that the intelligent driver model is not as accurate as other models when representing one driver, but when multiple drivers are considered, it can outperform some other models.

Table 21. IDM parameter values by driver.

	Driver J	Driver K	Driver F	Driver L
a	2.6519	3.3372	5.8950	0.8724

vdes	106.5559	97.3445	144.7275	96.0030
s0	27.5776	8.8639	2.5513	15.6821
T	0.0831	2.2913	2.3686	2.1768
b	1.0143	1.0593	67.6662	0.3113
RMSE	1.2260	1.2215	1.9211	1.0985

Table 22 presents the calibrated parameters for the velocity difference model of the four drivers. The values of the free-flow speed and desired velocity suggest that their use in this model will only cause error, and that a replacement with a calibration parameter would yield a more accurate representation of the actual behavior of drivers. The velocity difference model shows a root mean squared error that is similar and not too distant from that of the Wiedemann model. The velocity difference model also shows very stable behavior across the different drivers in Table 22. This shows that the velocity difference model can adequately represent the behavior of different drivers in a manner that exceeds or is on par with the other models.

Table 22. VDIFF parameter values by driver.

	Driver J	Driver K	Driver F	Driver L
FFS	48.5177	79.9246	123.4765	155.2547
Ln	5.0480	50.4894	36.6056	17.3944
Lag time	0.9870	3.5970	0.2989	2.4494
Sensitivity	12.7509	13.6386	1.0580	8.2774
Vdes	93.1806	106.0174	33.6502	76.6298
Form factor	4.6059	6.9404	0.9134	7.1708
l	8.2814	6.0162	6.6070	5.2652
tau	4.4006	3.4847	1.3662	2.9527
RMSE	1.1082	1.0848	1.1990	1.2454

Table 23 shows the calibrated parameters of the Wiedemann model for the four drivers. The drivers show different null acceleration values and different threshold parameters, which highlights that the drivers have different preferences for accelerating and decelerating while following. This also shows that the drivers make the decision to change accelerations at different points. The flexibility available in the Wiedemann model undoubtedly contributes to its performance in mimicking the behavior of real drivers. The Wiedemann model shows stable behavior across the different drivers, just like the velocity difference model, but with a small hiccup for Driver F. The intelligent driver model and the Gipps model showed this same hiccup for Driver F. Nevertheless, the Wiedemann model represents different drivers well, as shown in Table 23.

Table 23. Wiedemann parameter values by driver.

	Driver J	Driver K	Driver F	Driver L
Ln-1	4.4938	5.5533	4.5469	5.6797
AX	2.5254	5.0441	2.8730	5.6009
BX	3.0408	3.4054	3.4211	3.4349
EX	2.5139	3.0961	3.4273	2.6760
CX	40.0659	75.0421	62.4053	93.1305
CX2	39.0314	28.9785	73.2578	73.0260
CLDVCX	31.7622	56.7976	33.1589	51.0795
OPDV	-2.5116	-6.5796	-1.0677	-3.3071
bnull	0.1224	0.0747	0.2408	0.0716
bmaxmult	0.0991	0.1533	0.4271	0.3871
FaktorVmult	0.0389	0.1670	0.2002	0.1820
bminadd	-20.8984	-13.1674	-23.4076	-41.7209
bminmult	0.0334	0.0913	0.3202	0.3784
Vdes	42.6738	84.0379	88.7207	16.0671
FaktorV	0.9396	0.5216	0.5065	1.9585
RMSE	0.7541	1.0916	1.4735	1.0764

Table 24 presents the calibrated parameters of the Fritzsche model for the four drivers. The model shows similar null acceleration values for three of the drivers. Three of the drivers also show similar $an-1$ and $an+$ values. The Fritzsche model is very similar to the Wiedemann model in the way it operates, but the root mean squared error shows that the Fritzsche model does not perform similarly to the Wiedemann model. The Fritzsche model shows a competitive root mean squared error for Driver L, but it is lacking with the other drivers. This inconsistency shows in the Fritzsche model having mediocre performance compared to the other models for this particular set of data, which highlights its inadequacy for mimicking multiple drivers.

Table 24. Fritzsche parameter value by driver.

	Driver J	Driver K	Driver F	Driver L
Sn-1	22.6295	2.0673	2.0673	1.8332
kPTN	5.4213	4.8326	4.8326	4.0063
kPTP	1.1493	3.0874	3.0874	1.9313
fx	19.3430	-62.8779	-62.8779	-62.5547
Td	4.4129	4.8104	4.8104	3.9579
Tr	1.8111	0.8955	0.8955	1.0082
Ts	2.9147	3.6792	3.6792	3.8147
bmin	-2.4743	-0.0955	-0.0955	-0.8210
an-1	-4.5296	6.8066	6.8066	6.0010
bnull	-6.0129	6.7731	6.7731	6.4286
an+	-5.3233	-5.3827	-5.3827	-3.5638
RMSE	2.2094	1.3413	1.4332	1.1242

Figure 27 presents an example fitted car-following period for Driver L that is representative of the observations seen in other car-following periods for all four drivers. Most of the models show a behavior that is heavily influenced by the lead vehicle, as shown by the lines for the models overlapping the line for the lead vehicle in the figure. The actual behavior of the driver shows less dependence on the behavior of the lead vehicle, as shown by the black line in Figure 27. The GHR and Gipps models show large error in this car-following period. The error of the Gipps model is the constant velocity, while the error of the GHR model appears to be that this car-following period represents behavior that the calibrated GHR model does not adequately capture. The rest of the models show similar behavior that, as mentioned before, is heavily influenced by the actions of the lead vehicle.

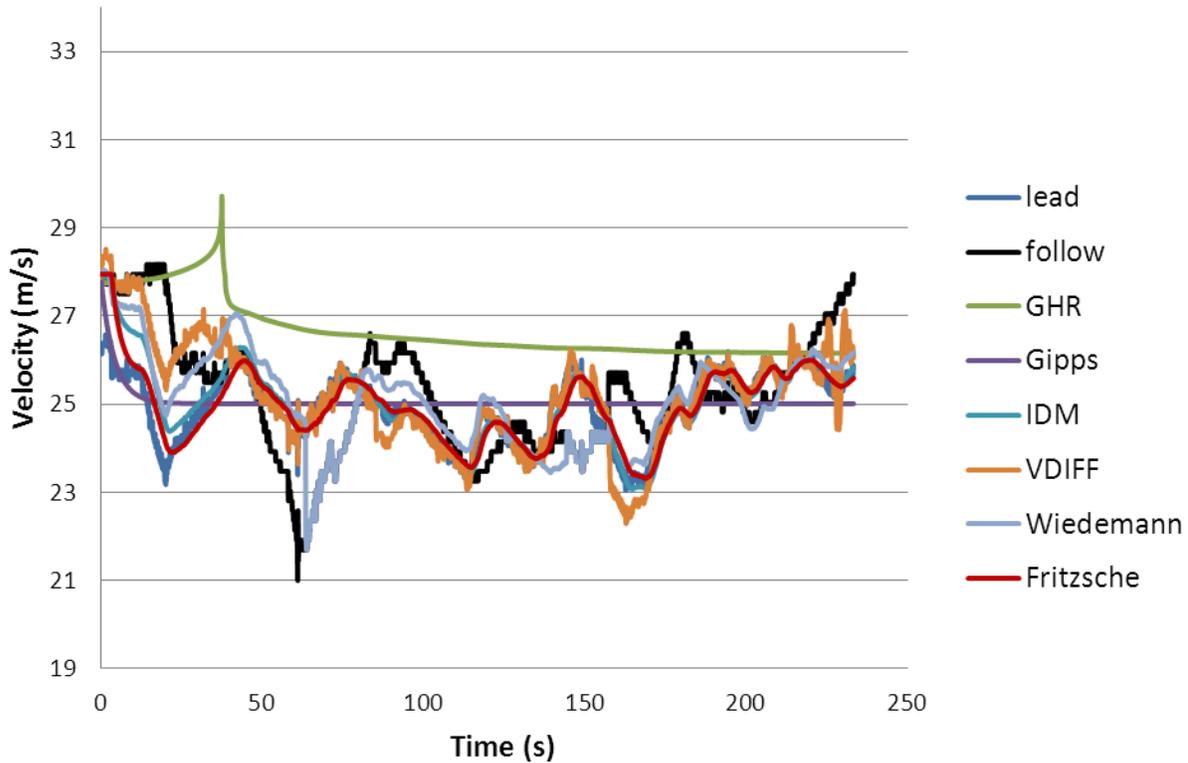


Figure 27. Graph. Plot of models as compared to the data for Driver L.

Conclusions

The results show that some of the microscopic traffic flow models calibrate to match real drivers better than others. The results also show that some of the models are more adequate at mimicking different truck drivers. Most of the models show a behavior that is heavily influenced by the actions of the lead vehicle. The results of this research effort suggest that the velocity difference model and the Wiedemann model can adequately represent the behavior of different drivers. This means that if a single car-following model is used with data from different drivers individually, these two models show the most promise in being up to the task. It is important to note that these findings are based solely upon data for truck drivers, so the applicability of these findings to car drivers' behavior would be questionable and thus is recommended as an area of future research.

CHAPTER CONCLUSION

The naturalistic data offers new insights into driver behavior due to the sheer amount of car-following data that is available for each driver. The amount of data allows for the discovery of new thresholds for the Wiedemann model: a passing threshold and a hook following threshold. These two new thresholds increase the usability and accuracy of the Wiedemann model when natural or naturalistic driver behaviors are concerned because these thresholds were developed based on observations made in the naturalistic data. Combining the strengths of both the GHR model and the Wiedemann model into a hybrid yields a more accurate and highly flexible model. The only downside to this combination is the introduction of additional calibration parameters,

which makes the model more difficult to calibrate. The comparison of different car-following models showed that different models have the least error for different drivers, but the velocity difference model and the Wiedemann model both perform adequately across all the drivers tested.

The hybrid model developed in this chapter allows better modeling of normal car-following behavior, and provides a better and seamless integration between car-following models of normal events and agent-based modeling of safety-critical events. The behavior of drivers during safety-critical events differ from their normal behavior and is better handled with the techniques presented in the following chapter because current models do not have the capacity to address the complex dynamics that occur during a safety-critical event.

CHAPTER 5. AGENT DEVELOPMENT AND TRAINING WITH REINFORCEMENT LEARNING

BACKGROUND INFORMATION

The goal of this research task is to model naturalistic vehicle actions during safety-critical events. This was to be achieved by simulating driving policy to associate observed traffic state to vehicle actions. Driving policy should be driver independent and consist of multiple driving rules that are associated with different states. Safety-critical events in this study are located in the emergency regime. Most car-following models, such as the Wiedemann model, assume that drivers will always keep a safe distance so that vehicles should always be able to stop before hitting lead vehicles in an emerging conflict. However, in real-world cases, driver actions are not always just longitudinal deceleration. For example, a driver may make a maneuver and execute a lane change. Also, as traffic stimuli and causalities sometimes vary case by case, it is very hard to establish predefined longitudinal and lateral action models. So far, no lateral action models have been used to simulate actions for safety-critical events.

TRAFFIC STATES AND ACTIONS DURING SAFETY-CRITICAL EVENTS

Drivers behave according to the traffic state they experience. Traffic state is defined by a set of variables that can represent a vehicle's kinematic conditions and its environment. In fact, in most safety-critical events, a vehicle's action is closely related to its interaction with a leading vehicle. As most car-following models suggest, vehicle speed, distance from the leading vehicle, and relative speeds are considered traffic state variables. Additionally, we assume driving during safety-critical events is a sequential task in which driver actions are continuous, so that earlier actions may contribute to current action decisions. For example, when a driver decides to decelerate in the previous state, he or she is more likely to continue decelerating in the current state instead of executing a maneuver. Therefore, actions of the previous state are considered to be variables of the current state.

AGENT-BASED MODELING OF DRIVER BEHAVIOR

Behavior rules of individual drivers should be used to teach an agent simulator (of an individual driver). Behavior rules associate actions with traffic state and provide a driver-dependent driving policy for its agent to follow. So when an agent experiences a certain traffic state, the policy will map the traffic state to associated actions. By using naturalistic driving data in training, agents will learn to adopt driving rules during the training procedure and should be capable of replicating driver and vehicle actions when training is completed. Accordingly, if data on safety-critical events is used in training, agents will perform driver-specific naturalistic actions, which could probably result in a crash or near-crash. Reinforcement learning can be used to determine the actions an intelligent agent should take in an environment to maximize some notion of long-term goals.⁽⁷⁴⁾ The objective of using reinforcement learning algorithms is to find a policy that maps traffic states to their optimal actions. In our research scope, the agent driving policy should be close enough to the "guiding" driver. One of the remarkable research papers by⁽⁷²⁾ illustrates the basic mechanism of reinforcement learning. Reinforcement learning reinforces agent actions when they perform approximately close to naturalistic actions and penalizes actions that are far

away from them. The only information available for learning is the system feedback, which describes the agent's task in terms of reward and punishment. At each time step, the agent receives a reinforcement signal regarding the last action it performed in the previous state. The challenge is optimizing not only the direct reinforcement, but also the total amount of reinforcements the agent can receive in the future. Finally, reinforcement learning should extract driving rules from a naturalistic dataset and establish similar driver-specific state action mapping rules.

POTENTIAL SHORTCOMINGS OF USING EXISTING REINFORCEMENT LEARNING ALGORITHMS FOR AGENT TRAINING

Although past research efforts have yielded great accomplishments, many limitations still exist. The dimension problem is a great difficulty of most agent controllers and simulators. If the number of traffic state variables is large, computation time would become a burden in training. Because the reinforcement learning mechanism has to associate state with action and update mapping rules, more dimensions in traffic state will cause an exponentially increasing number of mapping rules.

Moreover, the existing reinforcement learning efforts mainly deal with discrete state-action mapping problems so that mapping rules could refer to a look-up table to relate a discrete traffic state to a discrete action. However, traffic state variables (vehicle speed, for example), especially in our study, have continuous distributions, which makes it impossible to build a look-up table to relate states to actions.

Another significant drawback of the established reinforcement learning algorithm in transportation is the inability to generate continuous actions. Instead, most studies use discrete action sets consisting of a limited number of actions for the agent to select. Abdulhai's Q-learning mechanism has a binary action option phase: extension or end.⁽⁶¹⁾ Zhang improved this action selection methodology and enabled controller agents to select one of the six actions. This problem limits established reinforcement learning algorithms in continuous action simulation.⁽⁷⁵⁾

THE PROPOSED REINFORCEMENT LEARNING METHODOLOGY

In this research task, we proposed a revised reinforcement learning methodology to solve the traffic state dimensional problem and continuous action generation problem. The proposed revised neuro-fuzzy actor-critic reinforcement learning (NFACRL) approach can alleviate most of the computation burden that was associated with the previously developed approaches. NFACRL consist of an actor-critic reinforcement learning mechanism that performs a reinforcement learning training through two parts: (1) an actor that is responsible for producing an action corresponding to each state and (2) a critic that is responsible for calculating the long-term reward for the produced action and use the reward to reinforce good actions. The *neuro* part of the NFACRL mechanism uses a neural network approach to update all the weights responsible for mapping states to actions based on the actor's calculated reward. The *fuzzy* part of the NFACRL is responsible for transforming the continuous state space into single values representing degrees of membership in predefined discrete fuzzy sets, allowing the algorithm to handle high-dimensional state space. Using naturalistic traffic states and driving actions during crash and near-crash events, this approach can reproduce actual driver behavior during safety-critical events.

This chapter first describes the proposed revised NFACRL structure, including its framework as well as input and output components. Second, the naturalistic driving database and safety-critical events extraction process are described. Third, the agent training results are presented with a performance evaluation. Finally, cross validation between different agents is performed, driver heterogeneity is illustrated, and the idea of a mega-agent is proposed.

AGENT TRAINING METHODOLOGY: NFACRL

NFACRL Structure

NFACRL is a type of actor-critic reinforcement learning method. In this study, NFACRL uses current traffic state as inputs and generates vehicle actions as outputs. NFACRL then loads the state data of following states and decides the following actions. As its name indicates, NFACRL has a neural network structure and uses fuzzy logic to cluster traffic state variables and transfer information between layers. Accordingly, continuous traffic states are converted into a limited number of discrete fuzzy sets. Fuzzy rules are considered state-action mapping rules to associate fuzzy sets and actions. Reinforcement learning is the training process to obtain fuzzy rules.

As Figure 28 shows, our proposed NFACRL structure has four layers. The first layer is the input layer. Each of its nodes represents a state variable. The second layer is the fuzzy membership layer. Continuous state variables are fuzzified into linguistic terms such as “Speed is High” and “Speed is Low” in this layer. Each node is a fuzzy set with a membership function as output. Notice that one state input variable should have more than one fuzzy set so that continuous input variables can be transferred to discrete sets without losing much information. Fuzzy membership functions can be triangular, trapezoidal, and Gaussian depending on the distributions of state variables. The third layer is the fuzzy rules layer. Each rule is connected with a number of antecedents (discrete fuzzy sets) from fuzzy set layer. The firing strength is the output of a fuzzy rule and indicates the rule’s strength. The fourth layer is the actor-critic layer, which includes actor and critic nodes. Critic nodes are associated with the value of the next state after the agent chooses an action. Actor nodes are from one discrete action set. Each fuzzy rule selects one discrete action. NFACRL output action is the weighted average of all the actions selected by fuzzy rules in which fuzzy rule strengths are the weights. The following several paragraphs elaborate more on the framework design.

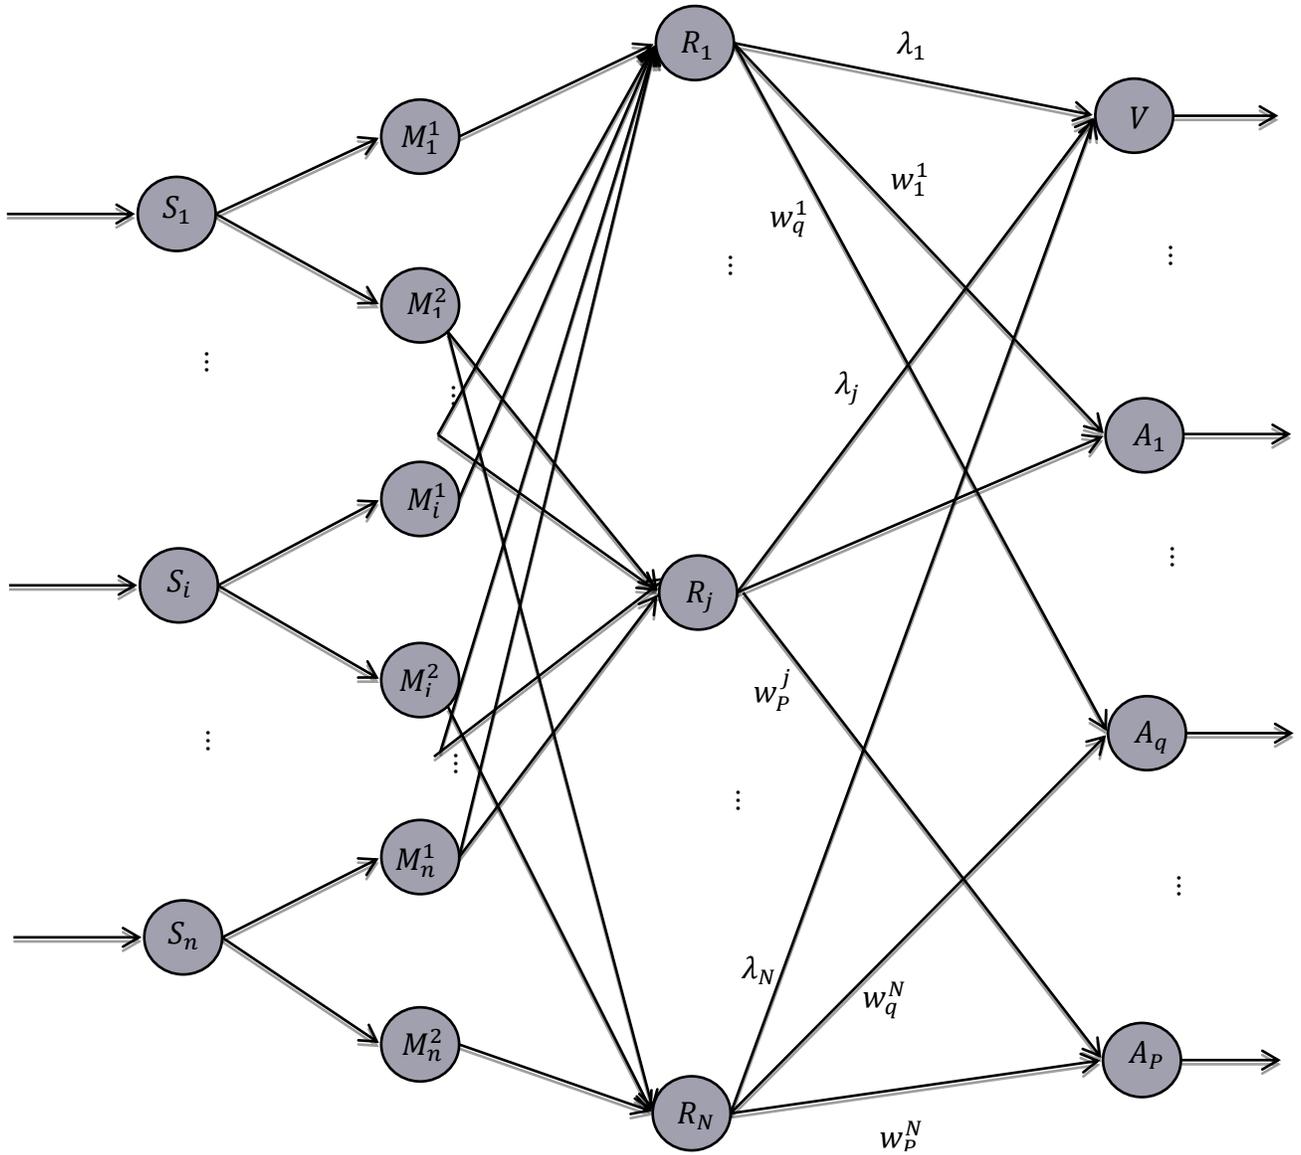


Figure 28. Diagram. NFA-CRL structure.

where

S_i is the i^{th} input variable (state).

n is the number of input variables.

NM_i is the number of fuzzy sets or membership functions for the S_i .

$M_i^{a(i)}$ is the $a(i)^{\text{th}}$ fuzzy set or membership function for the i^{th} input variable.

R_j is the j^{th} fuzzy rule.

N is the number of fuzzy rules.

λ_j is the weight between the j^{th} fuzzy rule and critic.

w_q^j is the weight between the j^{th} fuzzy rule and action q .

V is the critic output.

A_q is the output of the q^{th} discrete action.

$i = 1, \dots, n, a(i) = 1, \dots, NM_{i,j} = 1, \dots, m$ and $q = 1, \dots, P$.

Layer Design:

State Layer

Most safety-distance-based car-following models assume speed, relative speed, and distance as important stimuli when leading and following vehicles are interacting with each other. In addition, during safety-critical events, these three variables are also incorporated into state variables because the lateral lane offset (lateral position relative to the center of the lane), acceleration, and vehicle yaw angle (angle of vehicle longitudinal relative axis to lane markings) of the previous state are also important to the driver action decision process. Traffic state variables in the NFACRL structure are defined as in Equation 82 through Equation 87.

$$S_1 = v$$

Equation 82.

$$S_2 = \Delta x$$

Equation 83.

$$S_3 = \Delta v$$

Equation 84.

$$S_4 = \alpha'$$

Equation 85.

$$S_5 = y'$$

Equation 86.

$$S_6 = 0$$

Equation 87.

where S_i is the i^{th} state variable, v is the vehicle speed, Δx is the space headway (relative distance from the leading vehicle), Δv is the relative speed (speed of the leading vehicle minus the speed of the following vehicle), a' is the previous acceleration, y' is the previous yaw angle, and o is the vehicle lateral position offset.

Fuzzy Sets Layer

Each state variable in this study is linked to two fuzzy sets in the second layer. These two fuzzy sets are defined in linguistic terms: “ S_i is high” and “ S_i is low” for state variable S_i . Each fuzzy set is associated with a membership function. A triangular fuzzy membership function is used to make uniform a state variable and transfer its information to the two fuzzy sets. The triangular membership function requires upper and lower bounds of each state variable, which could be extracted from safety-critical events that the driver experienced. For each state, the lower bound and upper bound are the maximum and minimum state variable values. During the training process, upper and lower bounds are considered to be constant.

Membership functions for fuzzy sets Low and High are defined as:

$$\mu_{Low}(S_i) = \begin{cases} 1 & S_i \leq S_{lb,i} \\ \frac{S_{ub,i} - S_i}{S_{ub,i} - S_{lb,i}} & S_{lb,i} < S_i < S_{ub,i} \\ 0 & S_i \geq S_{ub,i} \end{cases} \quad i = 1, 2 \dots 6$$

Equation 88.

$$\mu_{High}(S_i) = \begin{cases} 0 & S_i \leq S_{lb,i} \\ \frac{S_i - S_{lb,i}}{S_{ub,i} - S_{lb,i}} & S_{lb,i} < S_i < S_{ub,i} \\ 1 & S_i \geq S_{ub,i} \end{cases} \quad i = 1, 2 \dots 6$$

Equation 89.

where $\mu_{Low}(S_i)$ is the membership function for fuzzy set “ S_i is Low,” $\mu_{High}(S_i)$ is the membership function for fuzzy set “ S_i is High,” $S_{lb,i}$ is the lower bound of state variable S_i , and $S_{ub,i}$ is the upper bound of state variable S_i .

Fuzzy Rules Layer

Fuzzy rules in the third layer provide a state action mapping policy to determine the actions to select from in the fourth layer. Each fuzzy rule is applied to one combination of fuzzy sets from the second layer. For example, a fuzzy rule can be represented as

WHEN “S₁ is low” and “S₂ is low” and “S₃ is high” and “S₄ is high” and “S₅ is high” and “S₆ is high,” THEN Action “Deceleration a=-0.05 g.”

In our designed neural network, each fuzzy rule is associated with six fuzzy sets originated from six continuous state variables. Because each state variable has two fuzzy sets, “Low” and “High,” the number of fuzzy set combinations should be $2^6 = 64$. Equivalently, there are 16 fuzzy rules in this study.

The product of the membership functions connected with each rule is used to compute its firing strength. Firing strength for the jth fuzzy rule is shown in Equation 90.

$$FS_{Rj} = \prod_{i=1}^6 \mu_{a(i)}^j (S_i)$$

Equation 90.

where a(i) is the linguistic term of the fuzzy set (either “Low” or “High”) for the ith input state variable, and j represents the jth fuzzy rule.

Actor and Critic Layer

Critic nodes represent the value of the next state under the current policy. Actor nodes represent actions to be selected by the fuzzy rules. Two types of actions, acceleration and yaw angle, are used as equivalents to vehicle longitudinal and lateral actions. Notice that acceleration and yaw angle are continuous variables. However, due to the constraints of neural network structure, it is impossible to enumerate all the possible continuous actions. Instead, discrete actions sets are defined here to include a number of representative action values for fuzzy rules to select. In this study, five discrete longitudinal acceleration values and five lateral yaw angle values are considered action candidates. Two sets are used to store actions, one for acceleration A and one for yaw angle Y (Equation 91 and Equation 92). Ten action values are considered as constant parameters during training.

$$A = \{a_{d1}, a_{d2}, a_{d3}, a_{d4}, a_{d5}\}$$

Equation 91.

$$Y = \{y_{d1}, y_{d2}, y_{d3}, y_{d4}, y_{d5}\}$$

Equation 92.

Neuron Weights

Weights are located between the fuzzy rule layer and the action layer. Critic weight λ^j links the j^{th} fuzzy rule to critic output. Action w_k^j links the j^{th} fuzzy rule to the k^{th} action output. Weights w_k^j show competition between discrete actions in the same action set and are used as references to the action selection methodology for fuzzy rule j . Weights λ^j are used to compute the value of traffic state V_s , which is the sum of the product of firing strength FS_{R_j} and weight λ^j (λ^j (Equation 93)).

$$V_s = \sum_{j=1}^{64} FS_{R_j} \lambda_j$$

Equation 93.

where FS_{R_j} is the firing strength for the j^{th} fuzzy rule R_j and V_s is the value of state S .

NFACRL Output Actions

Greedy Selection Algorithm

Fuzzy rule R_j selects discrete action α_{dk} from set A when weight w_{ak}^j is the largest among accelerations a_{d1} to a_{d5} . Similarly, R_j selects discrete action y_{dk} from set Y when weight w_{yk}^j is the largest in set Y .

Continuous Action Output

For each fuzzy rule, one discrete acceleration and one discrete yaw angle are selected. Therefore, the output actions acceleration a and yaw angle y are the weighted average of all the selected discrete actions by all the rules. Firing strengths of fuzzy rules are used as weights, so NFACRL agent actions are generated as seen in Equation 94 and Equation 95.

$$a = \sum_{j=1}^{64} FS_{R_j} * a_{dk}^j$$

Equation 94.

$$y = \sum_{j=1}^{64} FS_{R_j} * y_{dk}^j$$

Equation 95.

where a is continuous longitudinal acceleration and y is lateral yaw angle.

Instead of generating discrete actions, this NFACRL method is able to generate continuous action variables.

Weights Update: Reinforcement Learning Algorithm

The reinforcement learning algorithm updates the weights between the third layer and the fourth layer.

Weights Update

Weights λ^j and w_k^j are updated according to temporal difference (TD) error.⁽⁷⁴⁾ TD error consists of the value of the current state, the value of the following state, and a predefined reward function. The reward function is an estimation of how close the selected actions are at state S_t .

Temporal difference (TD) error is calculated as seen in Equation 96 and Equation 97.

$$\delta_{a,t} = r_{a,t+1} + \gamma V_{s_{t+1}} - V_{s_t}$$

Equation 96.

$$\delta_{y,t} = r_{y,t+1} + \gamma V_{s_{t+1}} - V_{s_t}$$

Equation 97.

where $r_{a,t+1}$ is the reward function when action a is taken at state S_t , $r_{q,t+1}$ is the reward function when action y is taken, γ is the discounting factor, V_{s_t} is the value of the current state, and $V_{s_{t+1}}$ is the value of the following state.

$V_{s_{t+1}}$ is calculated as seen in Equation 98.

$$V_{s_{t+1}} = \sum_{j=1}^N FS_{R_j,t+1} \lambda_{j,t}$$

Equation 98.

Notice that as the NFACRL agent generates two actions, two responding reward functions are computed and two temporal difference errors are calculated.

λ and w are updated by TD errors (Equation 99 through Equation 102), where

$$\lambda_{a,t+1}^j = \lambda_{a,t}^j + \beta \delta_{a,t} FS_{R_j,t}$$

Equation 99.

$$\lambda_{y,t+1}^j = \lambda_{y,t}^j + \beta \delta_{y,t} FS_{R_j,t}$$

Equation 100.

$$w_{ak,t+1}^j = w_{ak,t}^j + \beta \delta_{t,a} FS_{R_j,t}$$

Equation 101.

$$w_{yk,t+1}^j = w_{yk,t}^j + \beta \delta_{y,\alpha} FS_{R_j,t}$$

Equation 102.

where β is the learning rate.

TD errors update only the weights connected to the selected discrete actions a and y . From the equations in Equation 101 and Equation 102, when discrete action a_{dk}/y_{dk} is selected by rule R_j , only weight $w_{ak,t}^j/w_{yk,t}^j$ is updated.

Reward Function

The reward function provides guidance for an agent to follow. The reward function encourages an agent to take actions that are close to driver actions and penalizes actions that are dissimilar. The agent will know the outcome of the actions only when they have been chosen. When the performance of an action outcome is good, the reward function should be positive to provide a greater probability of the agent choosing those actions in the future. When performance is bad, the reward function should be negative to reduce the probability of the agent making those choices again. In this project, actions from the naturalistic driving database were used as references.

First, relative error is calculated as in Equation 103

$$e_a = \left| \frac{a - a_n}{a_n} \right|$$

Equation 103.

for acceleration and as in Equation 104

$$e_y = \left| \frac{y - y_n}{y_n} \right|$$

Equation 104.

for yaw angle, where e_a is the absolute relative error of acceleration, e_y is the absolute relative error of yaw angle, and a_n and y_n are the naturalistic actions from the database.

Then, a nonnegative parameter e_{th} is defined as an acceptance threshold. When e_a or e_y is less than e_{th} , the reward function (Equation 105 and Equation 106) is positive and related weights will increase, which increases the probability of fuzzy rules selecting a and y in the future.

$$r_a = \alpha(e_{th} - e_a)$$

Equation 105.

$$r_y = \alpha(e_{th} - e_y)$$

Equation 106.

where r_a is the reward function of acceleration, r_y is the reward function of yaw angle, and α is the scaling factor.

Initially, when an agent does not know which discrete action to take, it may choose an extremely inappropriate action, causing e_a and e_y to become too large. When that condition occurs, the reward function becomes negative enough to interrupt the weights update. Therefore, reward functions need to be adjusted to make sure they have a lower bound (Equation 107 and Equation 108).

$$r_a = -\alpha \quad \text{when } e_a \geq 1$$

Equation 107.

$$r_y = -\alpha \quad \text{when } e_y \geq 1$$

Equation 108.

NATURALISTIC DATA EXTRACTION AND AGENT TRAINING PROCEDURE

Database Description

In agent training, we use safety-critical events from the 8-Truck database of the Naturalistic Truck Driving Study (NTDS), the 34-Truck database, and the 100-Car database collected by VTTI. See Chapter 3 for further descriptions of these studies.

In our research focus, the following vehicle was the instrumented vehicle. The measured subject vehicle data included speed, yaw angle, lane offset, and accelerations. Range and range-rate were collected by instrumented forward viewing radar from the following vehicle. Speed was collected from the speedometer. Yaw angle and lane offset were extracted from video recording. Acceleration from the accelerometer was used as longitudinal traffic action, and yaw angle was used as lateral action.

Safety-Critical Events Extraction

The safety-critical events were identified and analyzed in previous work by VTTI. The method used to identify the safety-critical events were triggers or thresholds on individual variables that were collected.

The safety-critical events selected in this study are rear-end crash and near-crash conflict, especially when a following vehicle reacts to sudden braking or a sudden merge of the leading

vehicle from the adjacent lane. When an event happens, the driver of the following vehicle brakes or swerves to avoid the incoming conflict.

State and Action Variables Selection

Before training, the training parameters (upper and lower bounds of fuzzy sets and discrete action set values) were determined from the naturalistic data. In this experiment, we used minimum and maximum state variables as lower and upper bounds.

Because discrete action sets should represent naturalistic driving actions, one basic method of selecting state and action variables is to extract representative actions from the naturalistic driving data. Statistically, the five representative action values of accelerations and yaws can be the minimum, the lower quartile (25th percentile, cuts off the lowest 25 percent of data), the median (50th percentile, cuts off 50 percent of data), the upper quartile (75th percentile, cuts off the highest 25 percent of data), and the maximum. Notice that the agent output action is the weighted average of discrete actions. When using minimum and maximum discrete actions, the weighted average will make it impossible for the agent action to reach a naturalistic maximum and minimum. Therefore, the maximum value is multiplied by 1.2 when the maximum is positive and by 0.9 when the maximum is negative. The minimum value is multiplied by 0.9 when the minimum is positive and by 1.2 when the minimum is negative. This change extends the range of agent action.

Training Process

In our designed training process, at one time step of one event, fuzzy rules scan their associated weights, select the optimal actions, and update the weights. NFACRL continues updating the weights from the beginning of the events until the end. Then, NFACRL loads initial state variables from the same events. Theoretically, when the differences between critic and actor weights for two consecutive iterations become very small, the training process is considered to be finished. However, the proposed NFACRL is a heuristic methodology, so no global optimal agent behavior is guaranteed. Therefore, it is possible that the convergence of the weights may be premature and that NFACRL results in a local optimal solution. One way to avoid this premature convergence is to give the agent sufficient training iterations. In this event example, we tried 400 iterations in agent training. Because an average safety-critical event has about 1,500 timing steps (150 s), driver fuzzy rules have been trained 600,000 times (1,500*400), and the resulting agent should produce a near-optimal approximation of naturalistic driver behavior.

During the learning process, a memory discount factor γ , a learning factor β , and a reward function scaling factor α affect the learning speed of NFACRL. γ controls the memory fade speed, in which the value of recently occurring states are weighted more. β shows how fast the agent adjusts to the new information. α controls the magnitude of the reward function and weight change, and e_{th} controls the sign of the reward function. In this test, $\beta=0.6$, $\gamma=0.9$, $\alpha=10$, and $e_{th}=0.2$.

Based on our preliminary efforts, 10 truck drivers and 10 car drivers were selected. We used all the safety-critical events available from database. Therefore, training parameters for each driver should be global, especially fuzzy thresholds of the state variables and discrete actions in the two action sets A and Y. By using all safety-critical events, we attempted to enrich agent experiences and teach agents to react as drivers did. In addition, a greater number of events may prevent bias effects. Although conditions and causalities of different events can be totally different, NFACRL,

theoretically, should still capture rules of individual drivers quite well because different events are located in different state spaces dominated by different fuzzy rules.

AGENT TRAINING RESULTS

As an example, Figure 29 and Figure 30 show the agent's longitudinal action (acceleration) and lateral action (yaw angle) during one event of Driver Agent A. The blue scatter plot represents the naturalistic driving actions, and the green curves show the agent actions. NFACRL captures naturalistic driver behavior quite well in this event, with an R squared degrees of accuracy of 0.981 and 0.967 for acceleration and yaw angle, respectively. We also tested NFACRL performance on another event of the same driver (see Figure 31 and Figure 32).

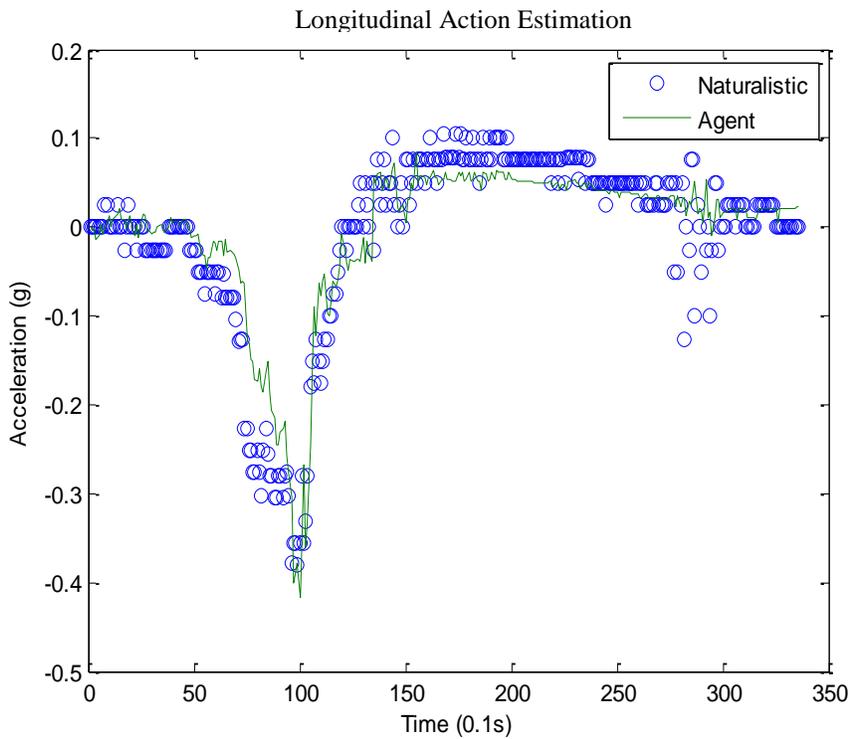


Figure 29. Graph. Acceleration of Agent A, event A1.

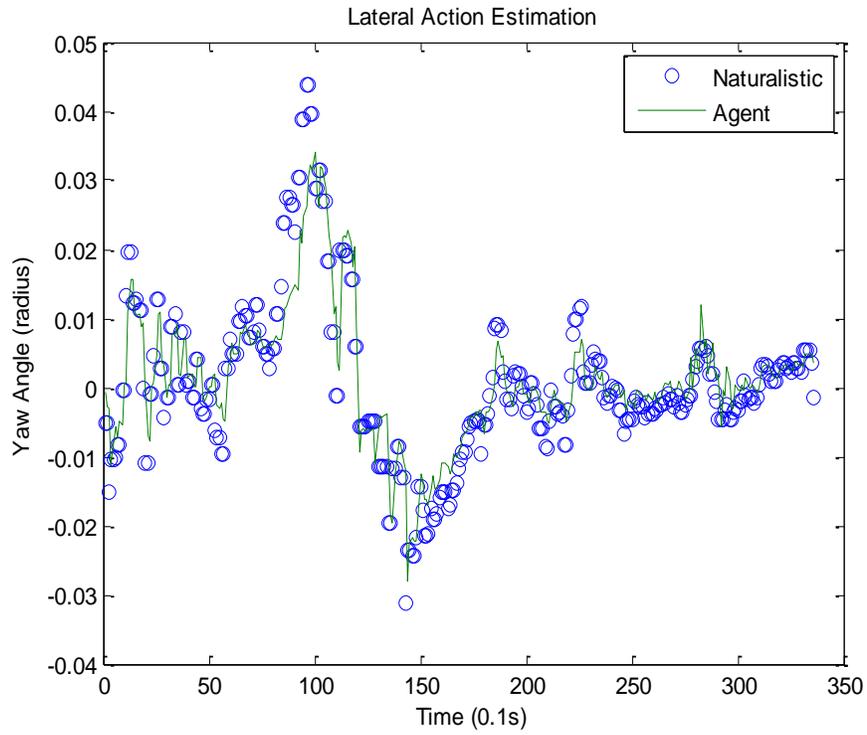


Figure 30. Graph. Yaw angle of Agent A, event A1.

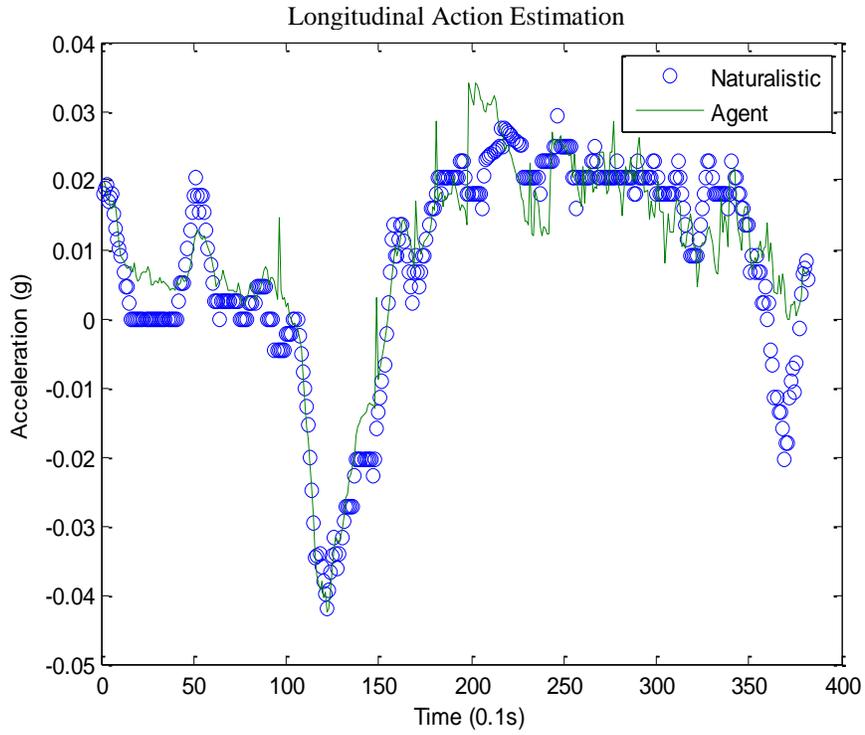


Figure 31. Graph. Acceleration of Agent A, event A2.

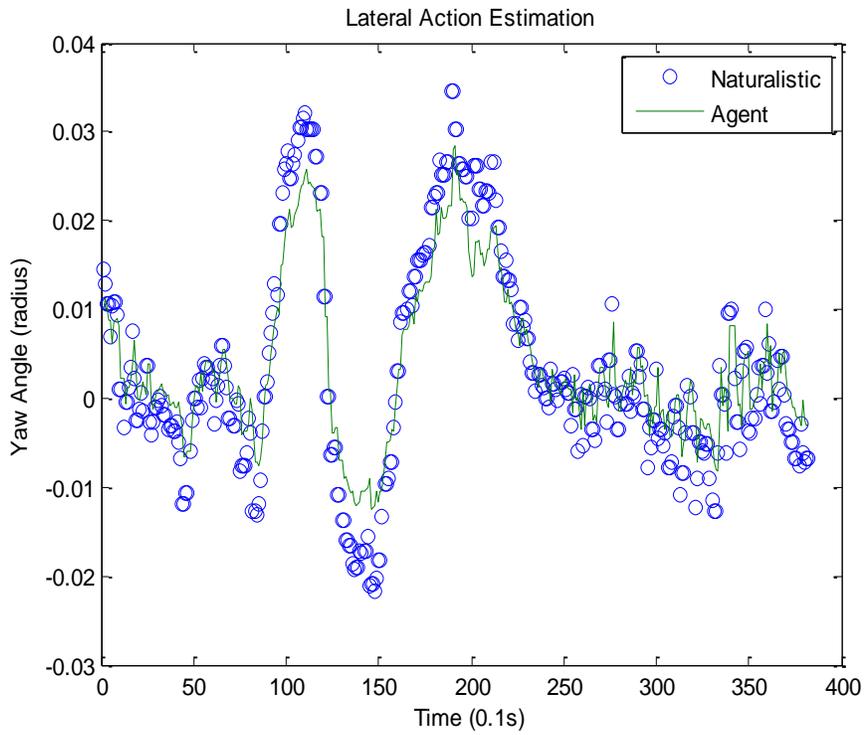


Figure 32. Graph. Yaw angle of Agent A, event A2.

Agent A captures Driver A's naturalistic behavior during events A1 and A2 quite well. It seems that even if events A1 and A2 vary, Driver A behaves similarly. This may prove that NFACRL methodology will not result in an average driving behavior, even when using more events as in training.

The four figures above show that during some part of the events, such as $t = 28$ s to 30 s in Figure 29, the action of Agent A diverges from naturalistic data by a small amount. Several reasons, including data collection errors in traffic state variables and unstable driver behavior within the events, may explain these differences. Occasionally, the leading vehicle does not fall into the range of the radar detection zone, so the agent assumes there is no vehicle in front of it. Consequently, the wrong traffic state leads to the wrong action. Human behavior is another phenomenon that is difficult to model. Within one event, a driver may react differently in similar traffic states. Also, the training factors may contribute to the error. From the reward functions, shown in Equation 107 and Equation 108, we tolerate NFACRL errors within a smaller amount of e_{th} , which may result in deviations from naturalistic actions.

CROSS VALIDATION RESULTS

The purpose of cross validation is to show the different behavior of different drivers when they experience the same hypothetical event. Driving behavior of one driver (for example, Driver A) must be obtained by using his or her events in training, and then states from another event (for example, event B) experienced by another driver (Driver B) should be used to demonstrate Driver A's actions during event B.

In our test, we trained Agent A and simulated Driver A's actions but using the event states from Driver B. Similarly, we simulated Agent B's actions using Driver A's events. Figure 33 and Figure 34 show an event example (event B1) from Agent B trained by using all its safety-critical events.

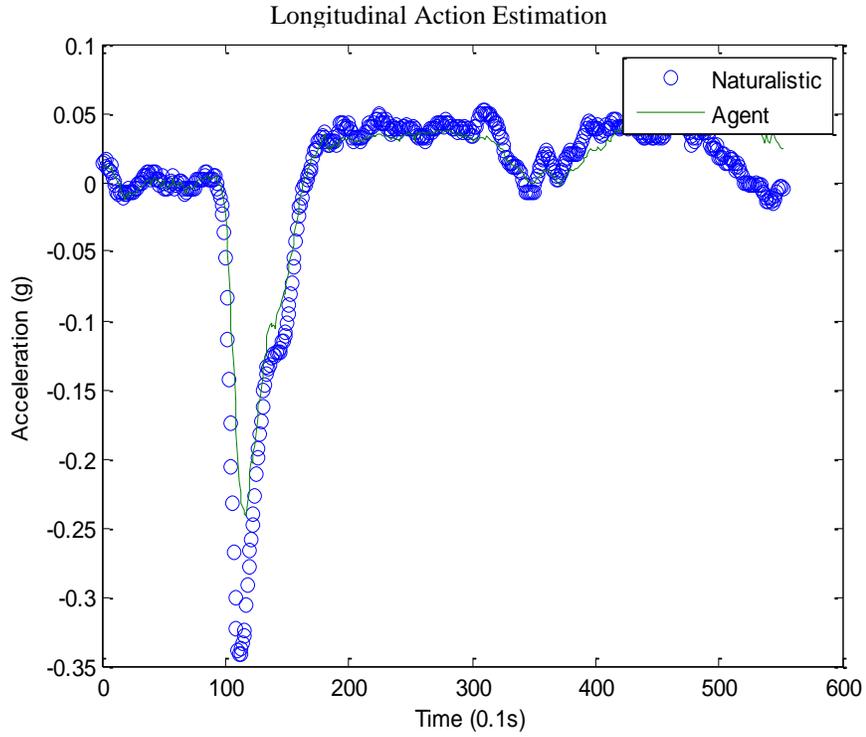


Figure 33. Graph. Acceleration of Agent B, event B1.

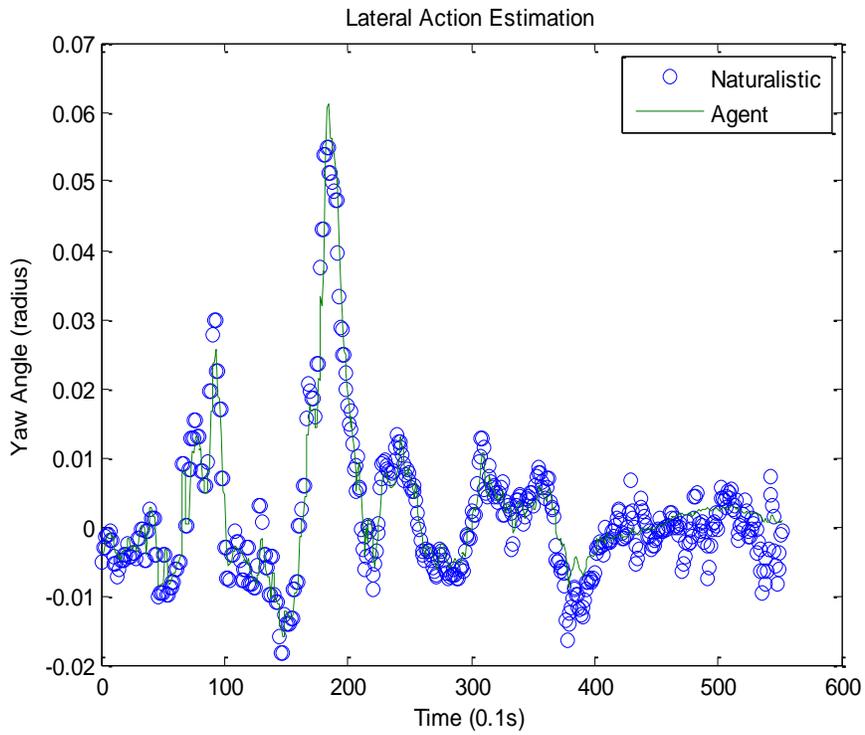


Figure 34. Graph. Yaw angle of Agent B, event B1.

Figure 35 and Figure 36 show the longitudinal and lateral actions of Agent B in one event from Driver A. Figure 37 and Figure 38 show the longitudinal and lateral actions of Agent A in one event from Driver B. It is very clear that Driver A and Driver B perform heterogeneous behaviors.

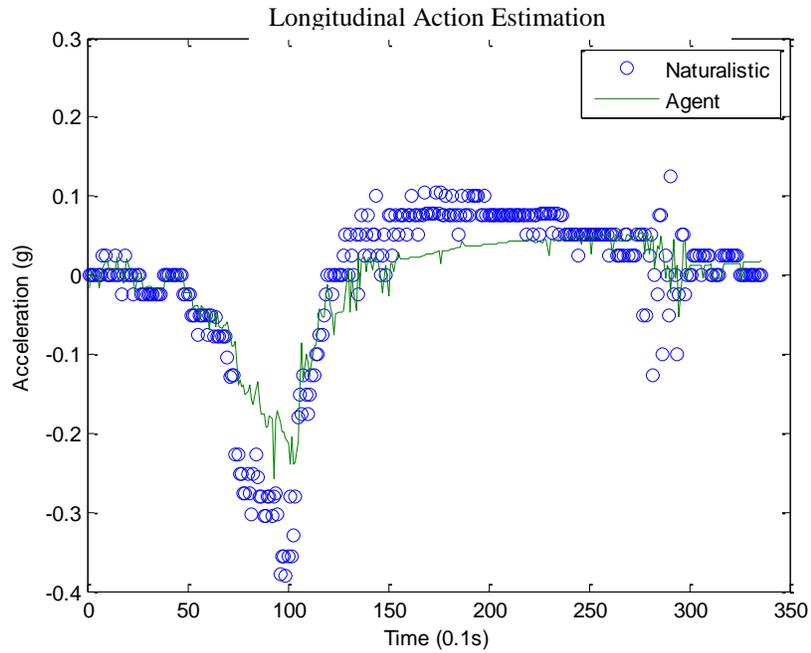


Figure 35. Graph. Acceleration of Agent B, event A1.

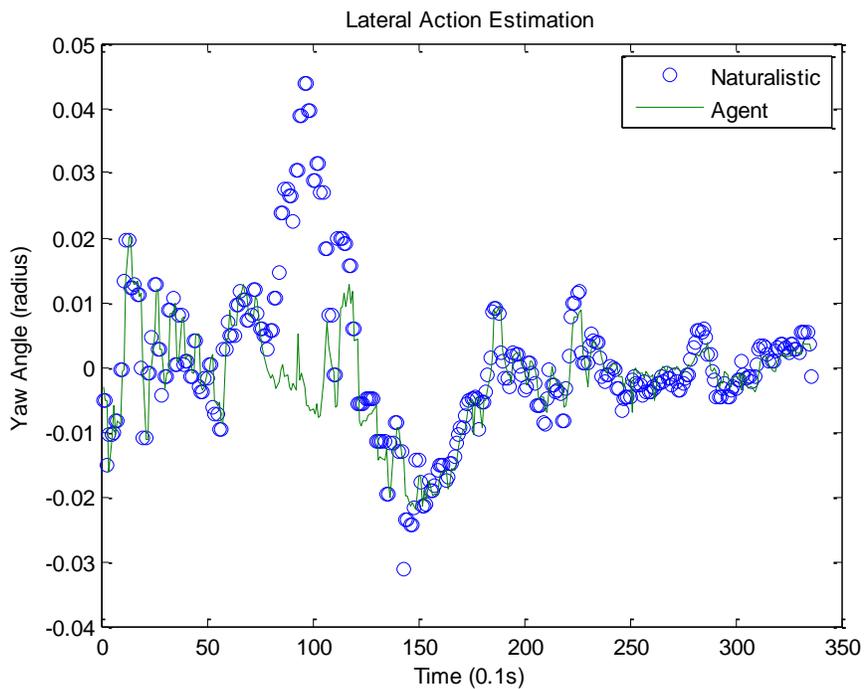


Figure 36. Graph. Yaw angle of Agent B, event A1.

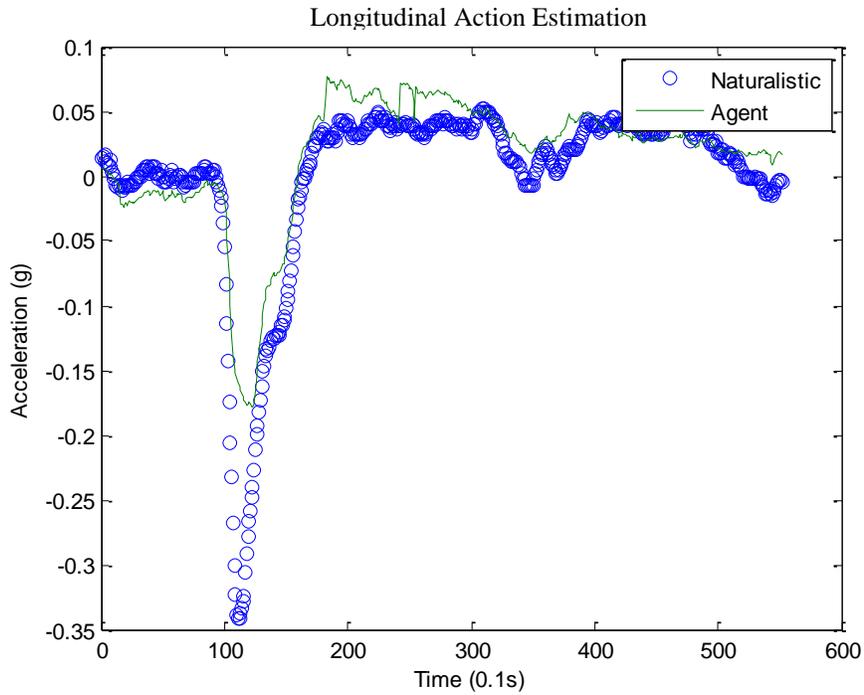


Figure 37. Graph. Acceleration of Agent A, event B1.

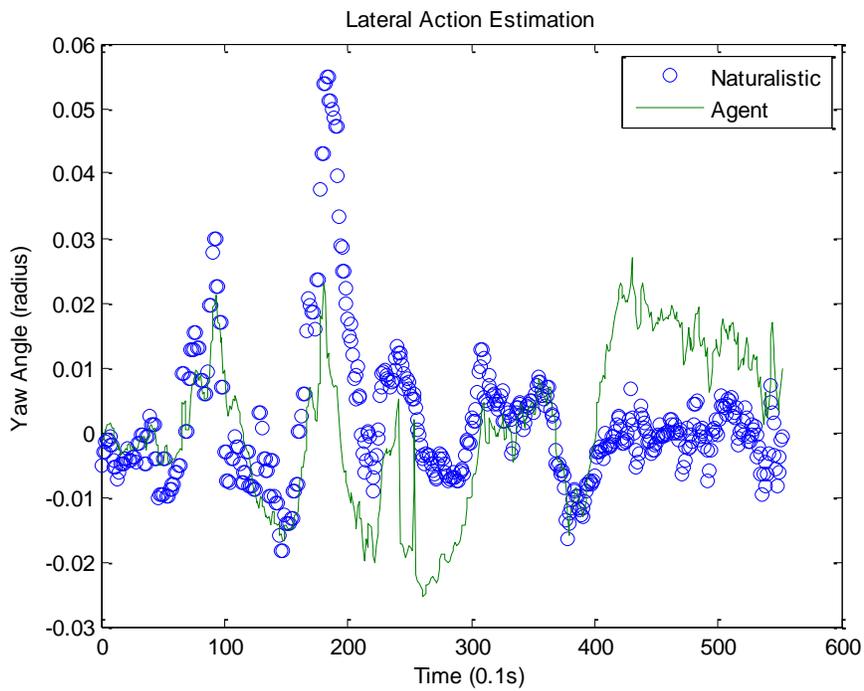


Figure 38. Graph. Acceleration of Agent A, event B1.

Table 25 uses R squared values as a statistical representation of the degree of accuracy of agent performance. In the table, the intersection of Driver A and Agent A and the intersection of Driver

B and Agent B represent the degree of approximation when using events of the same driver in training and validation. The table shows that the degree of approximation when using events of the same driver in training and validation is much greater than the degree of approximation in the cross validation (upper right and lower left), which also shows clear heterogeneities between the two drivers.

Table 25. R squared values for cross validation.

Event	Agent A Acceleration	Agent A Yaw	Agent B Acceleration	Agent B Yaw
Driver A	0.98	0.97	0.81	0.83
Driver B	0.82	0.60	0.97	0.92

THE MEGA-AGENT CONCEPT

We attempted to design an imaginary agent that can capture the behaviors of both Driver A and Driver B but would not result in an average behavior between A and B. Our revised NFACRL methodology may handle this challenge. According to the nature of traffic state variables, the state space in this problem has six dimensions. Differences in one dimension (lower speed and high speed, for example) can make two separate states in the state space. From the naturalistic data, as we found out, state variables from different events vary greatly, so it is less likely that two events would have many overlaps in the state space. Notice that in NFACRL, different fuzzy rules dominate different regimes of the state space divided by fuzzy sets of state variables. Therefore, the training process of the hypothetical mega-agent involves adjusting the fuzzy rules that are taking charge of relevant state-space regimes in which safety-critical events are located. Similarly, the fuzzy set thresholds and discrete action sets of training parameters are adjusted.

We used all the events of Driver A and Driver B to train the mega-agent. Performances of the mega-agent (using events A1 and B1) are presented in Figure 39, Figure 40, Figure 41, and Figure 42.

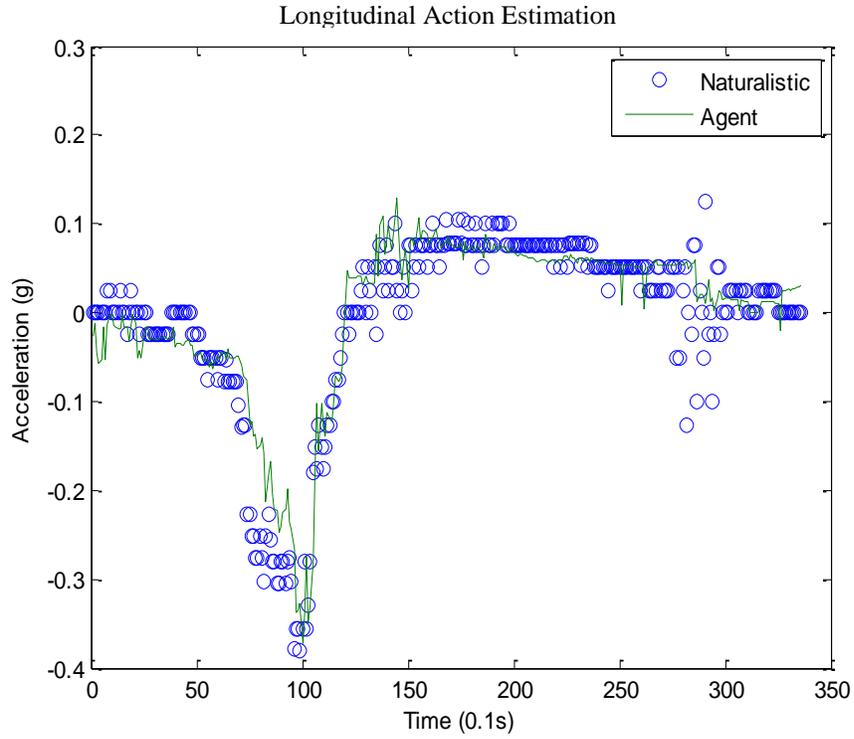


Figure 39. Graph. Acceleration of mega-agent, event A1.

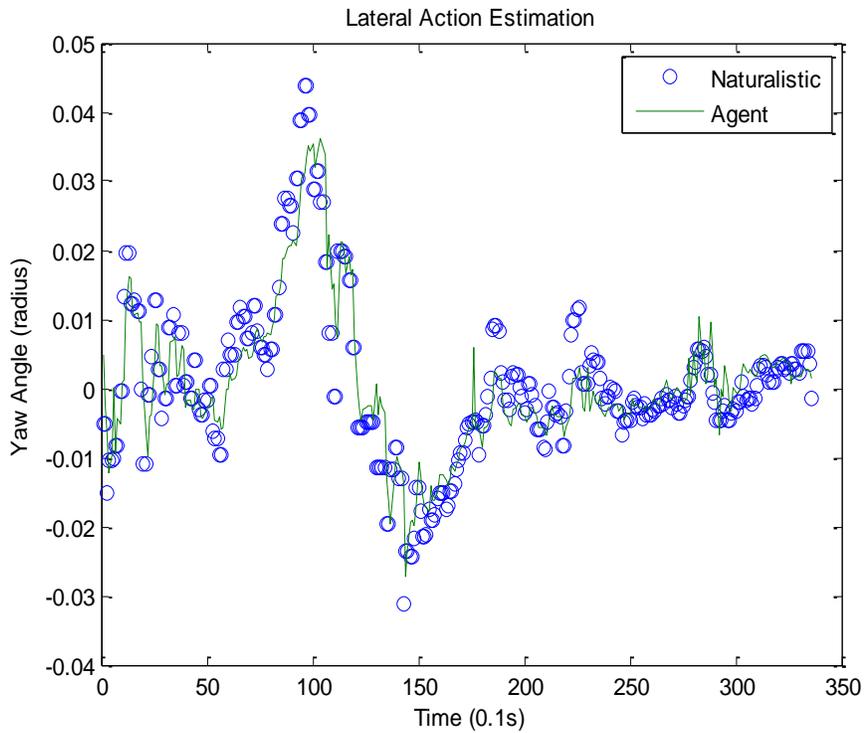


Figure 40. Graph. Yaw angle of mega-agent, event A1.

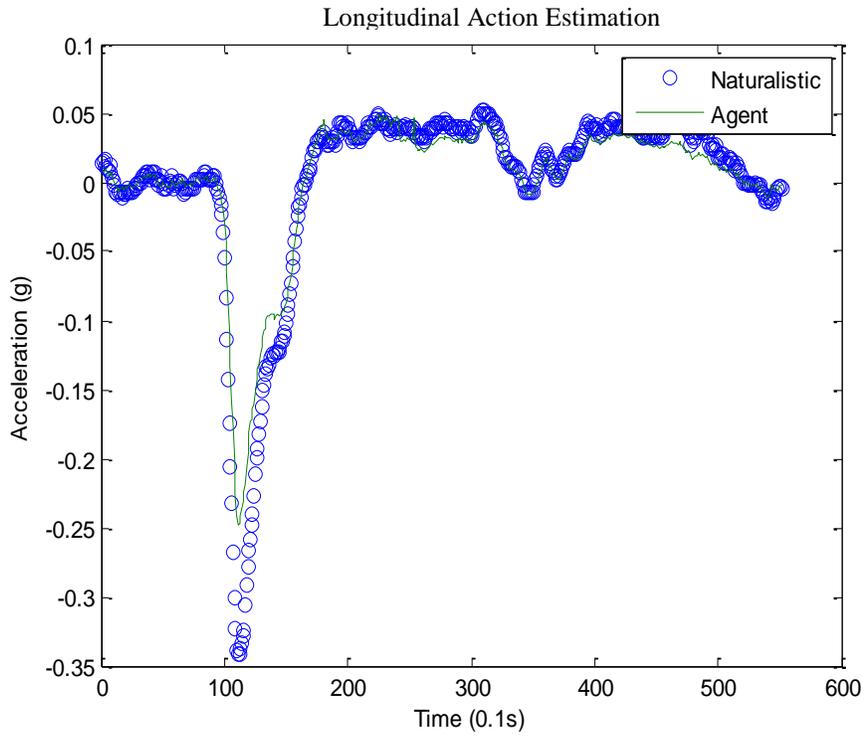


Figure 41. Graph. Acceleration of mega-agent, event B1.

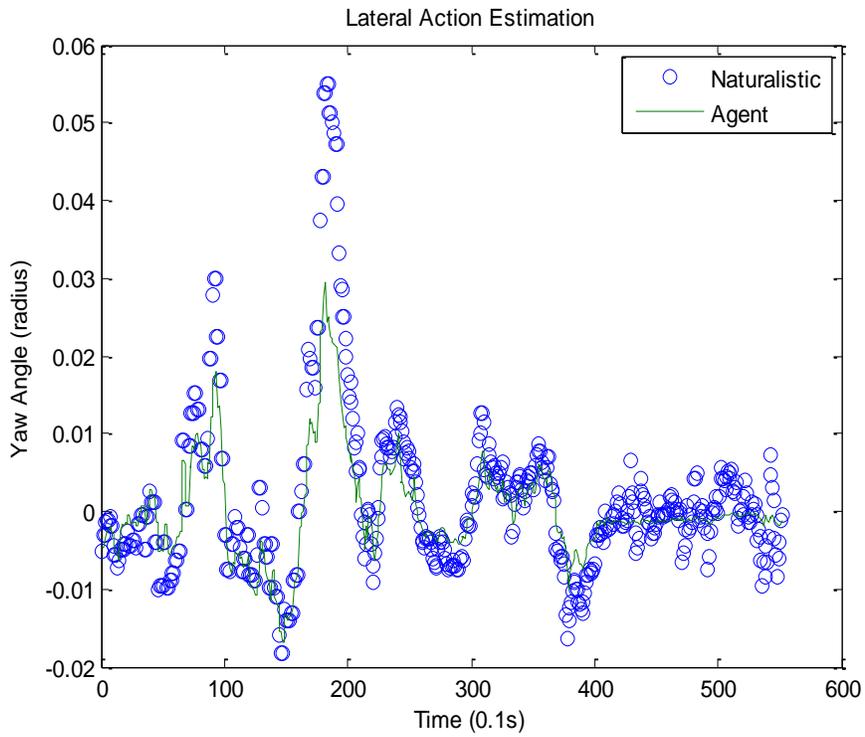


Figure 42. Graph. Yaw angle of mega-agent, event B1.

Compared to Figure 35, Figure 36, Figure 37, and Figure 38, the mega-agent performs better in cross-validation and can differentiate between Agent A and Agent B most of the time.

In reality, a conservative driver might never experience safety-critical events, but an aggressive driver may experience many. In that case, the conservative driver has no idea what actions to take to get out of a sudden emergency regime. Through this mega-agent training, the conservative driver will learn the crash avoidance actions from the aggressive driver and be better equipped to evade upcoming crashes.

Table 26 shows the R squared values of the mega-agent. They are quite high, although they are slightly lower than those when training and validation use the same events. In fact, the mega-agent is capable of mimicking the behaviors of Driver A and B at the same time without losing much driver specificity.

Table 26. R squared values of the test mega-agent.

Event	Agent A		Agent B		Mega-Agent Acceleration	Mega-Agent Yaw
	Acceleration	Agent A Yaw	Acceleration	Agent B Yaw		
Driver A	0.98	0.97	0.81	0.83	0.98	0.95
Driver B	0.82	0.60	0.97	0.92	0.97	0.91

From our preliminary efforts, we designed two mega-agents: one truck mega-agent and one car mega-agent. We selected safety-critical events from 10 truck drivers and 10 cars drivers to train the two agents separately. Each of them used all the event data from their 10 selected drivers in training. In fact, each of the two agents should know all the evasive behaviors of their 10 guiding drivers.

Table 27 shows the performance of the truck mega-agent. Two first two columns (under “Agent”) show the optimal performance when training data and validation events come from the same driver. As expected, compared to the agent itself, the performance of the mega-agent is approximately the same, although it is slightly worse in some of the cases because of the noise.

Table 27. R squared values of the truck mega-agent.

Event	Agent		Mega-agent		Cross validation with Agent A	Cross validation with Agent A
	longitudinal acceleration	Agent lateral acceleration	longitudinal acceleration	lateral acceleration	longitudinal acceleration	lateral acceleration
Driver A	0.97	0.97	0.98	0.97	0.98	0.97
Driver B	0.97	0.94	0.97	0.91	0.82	0.60
Driver C	0.98	0.96	0.97	0.96	0.93	0.86
Driver D	0.99	0.92	0.99	0.88	0.86	0.64
Driver E	0.88	0.96	0.81	0.95	0.47	0.76
Driver F	0.98	0.96	0.94	0.96	0.83	0.43
Driver G	0.86	0.98	0.84	0.88	0.86	0.62
Driver H	0.96	0.99	0.95	0.99	0.63	0.98
Driver I	0.95	0.98	0.93	0.98	0.48	0.75

Driver J	0.85	0.98	0.85	0.97	0.66	0.32
----------	------	------	------	------	------	------

Similarly, the degree of accuracy of the car mega-agent is presented in Table 28.

Table 28. R squared valued of the car mega-agent.

Event	Agent longitudinal acceleration	Agent lateral acceleration	Mega-agent longitudinal acceleration	Mega-agent lateral acceleration	Cross validation with Agent D longitudinal acceleration	Cross validation with Agent D lateral acceleration
Driver a	0.94	0.94	0.93	0.70	0.90	0.50
Driver b	0.97	0.94	0.96	0.79	0.82	0.30
Driver c	0.93	0.97	0.97	0.83	0.93	0.77
Driver d	0.97	0.93	0.96	0.93	0.97	0.93
Driver e	0.95	0.98	0.91	0.98	0.68	0.95
Driver f	0.97	0.92	0.96	0.89	0.92	0.86
Driver g	0.98	0.93	0.99	0.88	0.98	0.79
Driver h	0.98	0.93	0.99	0.89	0.98	0.90
Driver i	0.97	0.95	0.98	0.91	0.90	0.85
Driver j	0.95	0.92	0.99	0.89	0.98	0.65

CHAPTER CONCLUSION

This chapter reported a framework that simulates individual driver actions during safety-critical events. Neuro-fuzzy actor-critic reinforcement learning (NFACRL), an agent-based artificial intelligence machine-learning technique, was used to model driving behavior. The naturalistic driving database was used to train and validate driver agents.

The advantages of NFACRL lie mainly in its ability to simulate heterogeneous vehicle actions in complicated traffic environments. It is worth mentioning that this research was an attempt to apply reinforcement learning techniques to solving high-dimensional state problems and continuous-action simulation problems simultaneously in transportation research, especially in traffic flow theory. From the perspective of microscopic traffic behavior modeling, the proposed methodology is able to simulate lateral action, which brings new insights to the modeling of driver maneuvering behavior during safety-critical events. The proposed methodology also simulated events from different drivers and proved behavior heterogeneities.

This task focused on safety-critical events. For the technical part of this task, the training parameters speed, memory, and scaling factor were fixed for all the events, but it would be interesting to see how NFACRL performs under different combination of these factor sets. As with most artificial intelligence methods, NFACRL is a heuristic approach, so there is no guarantee that training results would be optimal. Therefore, training parameters play an important role in guiding NFACRL to find a near optimal result. During our test, we found that agent performance is very sensitive to the driver-dependent training parameters, discrete action sets, and state bounds. Theoretically, it would be better if these parameters were preoptimized before training. However, because the only way to test agent performance is through NFACRL training, parameter optimization and training form a cycle, and computation time would exponentially increase. Our approach set the parameters based on statistical quartiles, and this yielded good

results. Agents have the ability to handle lateral movements, which is outside of the abilities of most car-following models, but the ability to accurately judge when a safety-critical event occurs would be beneficial for knowing when the agent is crucial for an accurate simulation. The next chapter covers a discriminant analysis of normal driving behavior and safety-critical event behavior in order to find a way to distinguish between the two groups.

CHAPTER 6. SAFETY-RELATED AGENT ACTIVATION

DISCRIMINANT ANALYSIS

This research task focuses on using a statistical method, discriminant analysis, to find an accurate way to distinguish between normal car-following behavior and safety-critical event behavior. Discriminant analysis is a statistical method that uses information known *a priori* about the classification of observations from different datasets to distinguish between future observations from the same datasets (in this case two datasets: safety-critical and normal). The technique works by finding coefficients for the input variables that when multiplied by the variables and summed create a value that can be reliably used to classify future variables. In this chapter, we use the discriminant analysis technique to distinguish between safe (normal) and unsafe (safety-critical) driving patterns before a safety-critical event occurs. This is a concept that we present as a potential trigger for agent activation. In other words, if agent-based simulation were to be activated only to handle the simulation of safety-critical events, then the agents should be activated just before a safety-critical event is expected. The trigger-based activation could also be set with values that can ensure the activation of agents during the whole simulation period if it was so desired.

For this analysis, 30 datapoints from different car-following periods were used for each driver, along with the event data for that driver. This number of car-following or normal-driving points was used to produce a fair representation of normal driving behavior, while not overpowering the safety-critical event data in the analysis. Six variables were used for the discriminant analysis: longitudinal acceleration, vehicle speed, yaw angle, lane offset, range, and range rate. Equation 109 shows the mathematical form of the resulting discriminant score and how it relates to the coefficients for each variable.

$$\text{Discriminant Score} = \sum \beta_i * X_i$$

Equation 109.

where β_i is the coefficient for variable i and X_i is variable i .

Car Drivers

Table 29 presents the coefficient values resulting from a discriminant analysis at different time steps before the occurrence of a safety-critical event. The table shows that consistency exists in some of the coefficients over time, while the rest of the coefficients appear to be time dependent. Figure 43 shows the percentage of misclassification error at each time step. The time step at 6 s is chosen because it is the greatest amount of time that still maintains relatively low error. Figure 44 shows the results of applying the coefficients of the different time steps to the near-crash data as well as the normal car-following data. The results show some separation between the datasets at each time step, with some overlap occurring. Figure 44 shows that the error weighs on the side of misclassifying safety-critical events as normal driving behavior, which is not an acceptable error.

Time step 6 shows a clear separation between the two datasets, which means that the coefficients belonging to that time step would be the optimal choice for use in further study.

Table 29. Discriminant coefficients for Driver 1.

Time	Longitudinal Acceleration	Lateral Acceleration	Speed	Lane Offset	Yaw Angle	Range	Range Rate
0	5.4	2.7	0.051	0.019	-0.4	0.013	0.12
1	1.5	-13.1	0.013	0.019	1.6	0.013	0.15
2	-10.3	-10.6	0.030	0.022	-5.4	0.030	0.18
3	-4.0	-5.9	0.057	0.017	2.1	0.024	0.19
4	1.8	-8.7	0.053	0.014	15.3	0.000	0.13
5	-17.9	14.2	-0.005	-0.003	-16.3	-0.004	-0.15
6	16.5	-28.1	-0.026	0.014	38.0	0.001	0.15
7	9.8	-12.2	0.006	0.009	23.2	0.014	0.16
8	22.0	-17.5	-0.009	0.016	37.5	0.008	0.20
9	-14.8	15.3	0.006	-0.007	-3.0	-0.004	-0.15
10	14.3	-12.3	-0.004	0.014	-3.5	0.007	0.14

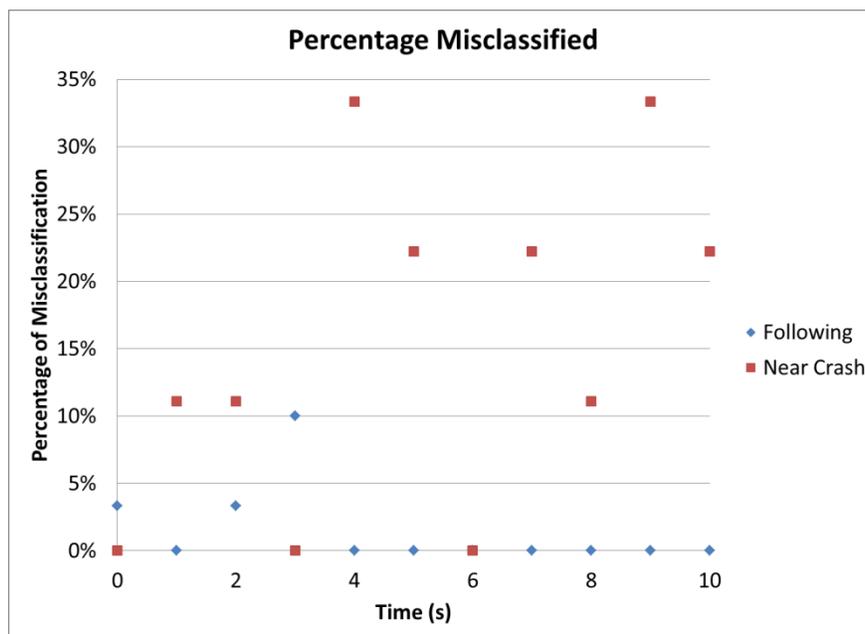


Figure 43. Graph. Percentage of misclassified datapoints for Driver 1.

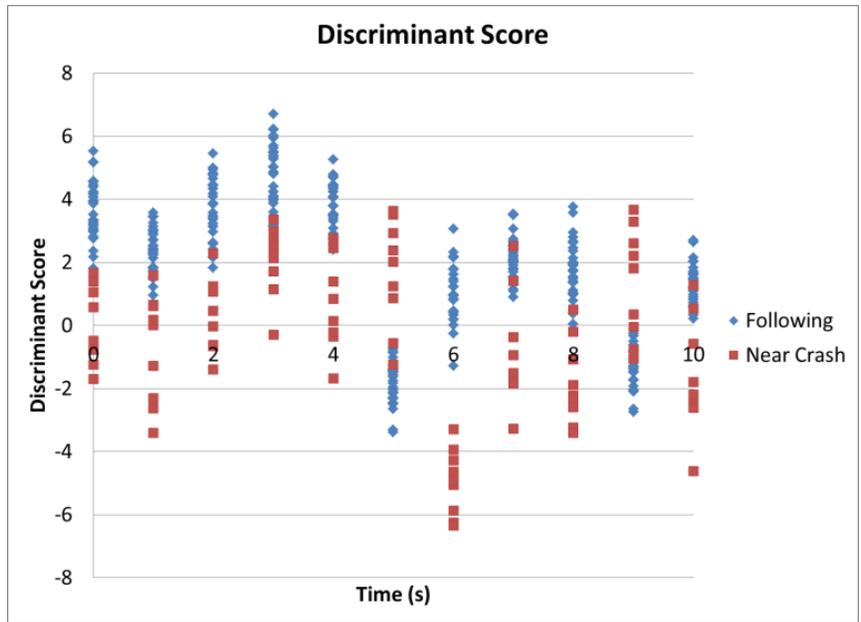


Figure 44. Graph. Discriminant scores for Driver 1.

Figure 45, Figure 46, and Figure 47 show the graphs of some of the terms that constitute the discriminant score for Driver 1. These graphs are the result of multiplying the input variables by the corresponding discriminant coefficient. The graphs show that the observable results of the discriminant analysis are not due to any one variable. Rather, the results are highly reliant on the specific combination of all of the variables that the discriminant coefficients create.

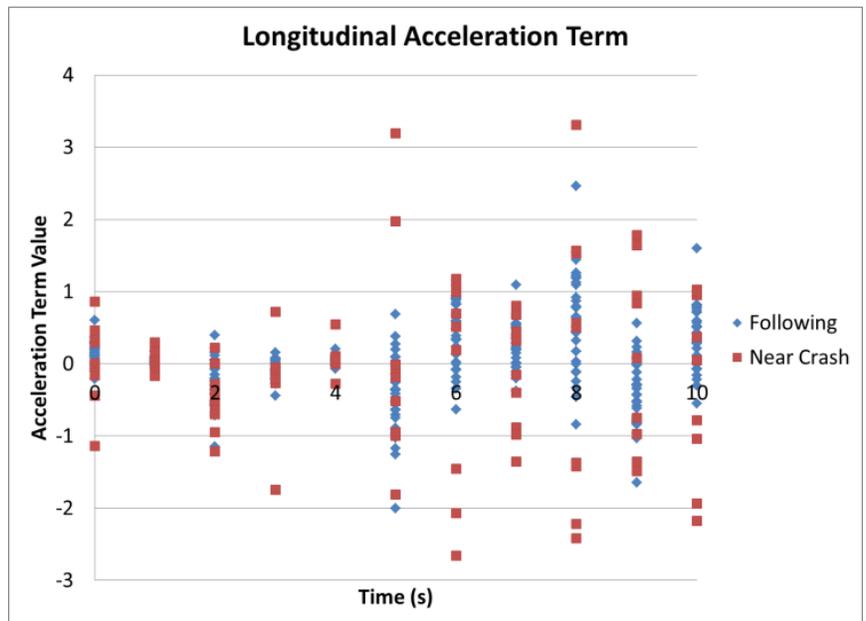


Figure 45. Graph. Longitudinal acceleration term for Driver 1.

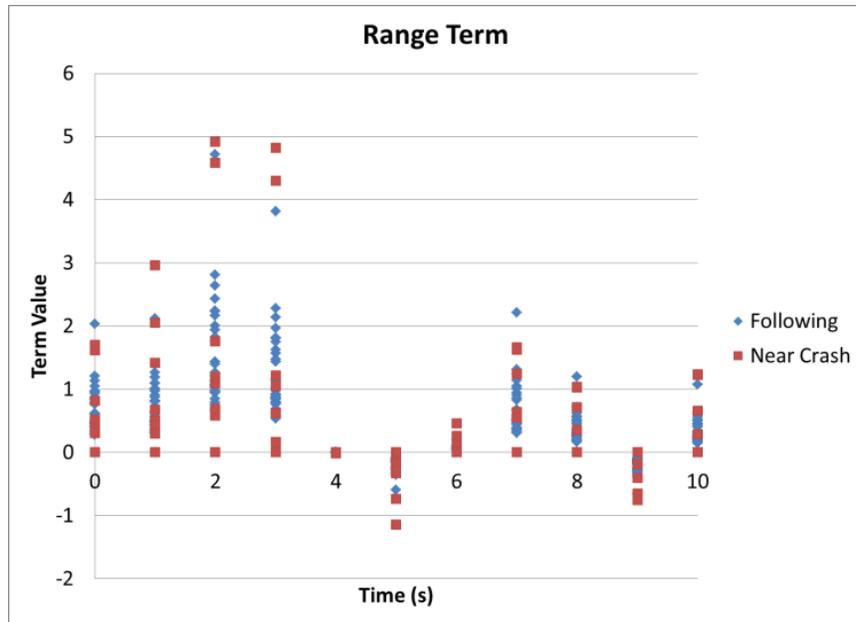


Figure 46. Graph. Range term for Driver 1.

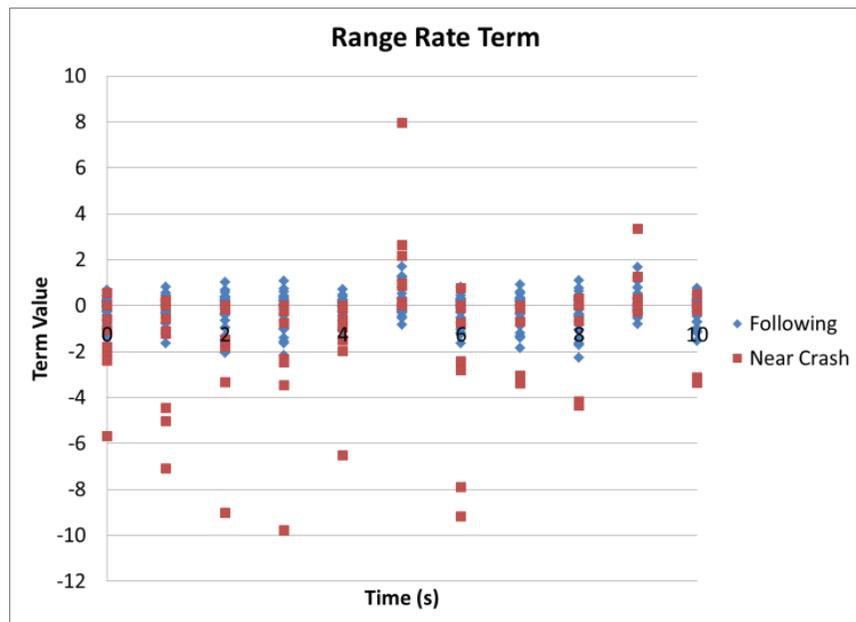


Figure 47. Graph. Range rate term for Driver 1.

Figure 48 shows the resulting discriminant scores when the coefficients from the 6-s time step are applied to a full car-following period. The figure shows some erratic behavior, but the discriminant scores are still above -3 , which is the highest point of the near-crash data in Figure 48. Figure 49 shows the results of applying the same coefficients to a sample near-crash event. The discriminant scores seem similar to a normal car-following period until the score suddenly

jumps below -5 . This jump corresponds to a vehicle suddenly changing lanes to be in front of the subject vehicle.

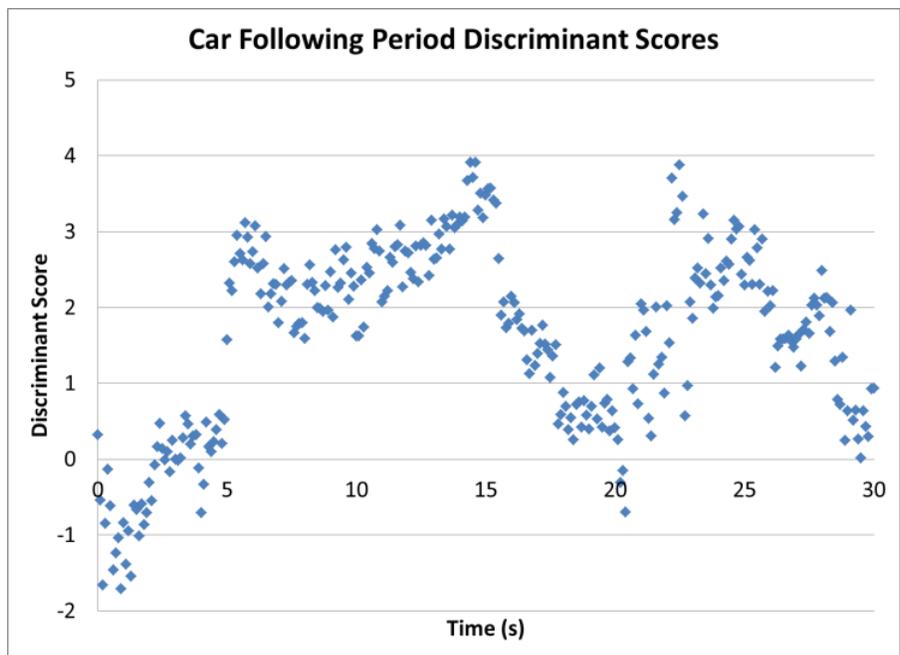


Figure 48. Graph. Discriminant scores of a sample car-following period.

Figure 50 shows the speed trajectory of the same near-crash event along with a marker of the 6-s prediction given by the discriminant coefficients. The 6-s prediction is well before the driver's reaction. Also, the discriminant score plateaus around -5 immediately after the 6-s prediction, which means that the driving conditions remain dangerous, but the driver has not adjusted to the change in conditions. When the driver decelerates, as seen in the trajectory in Figure 50, the discriminant score increases, indicating a change from dangerous conditions to safe conditions.

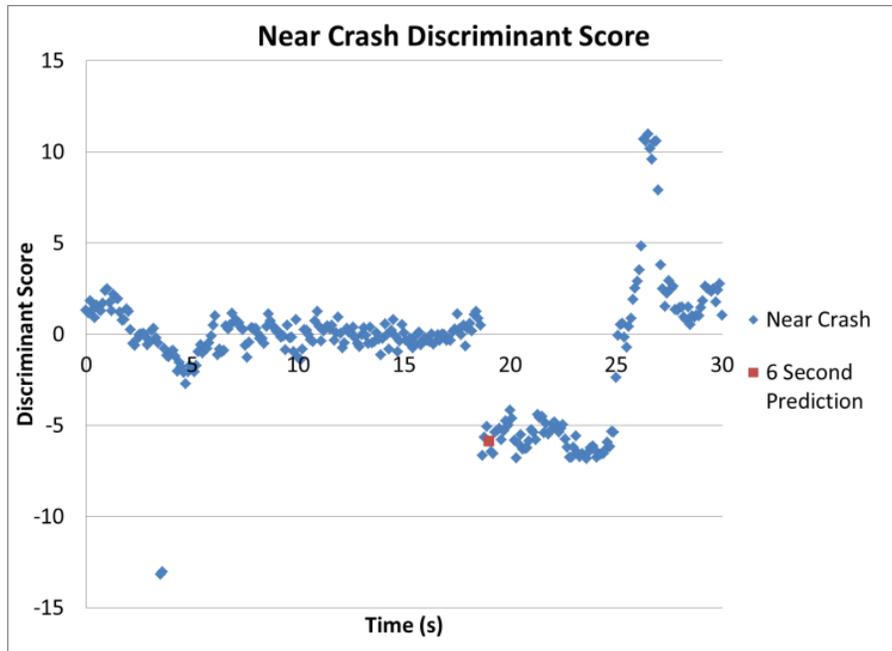


Figure 49. Graph. Discriminant scores of a sample near-crash event.

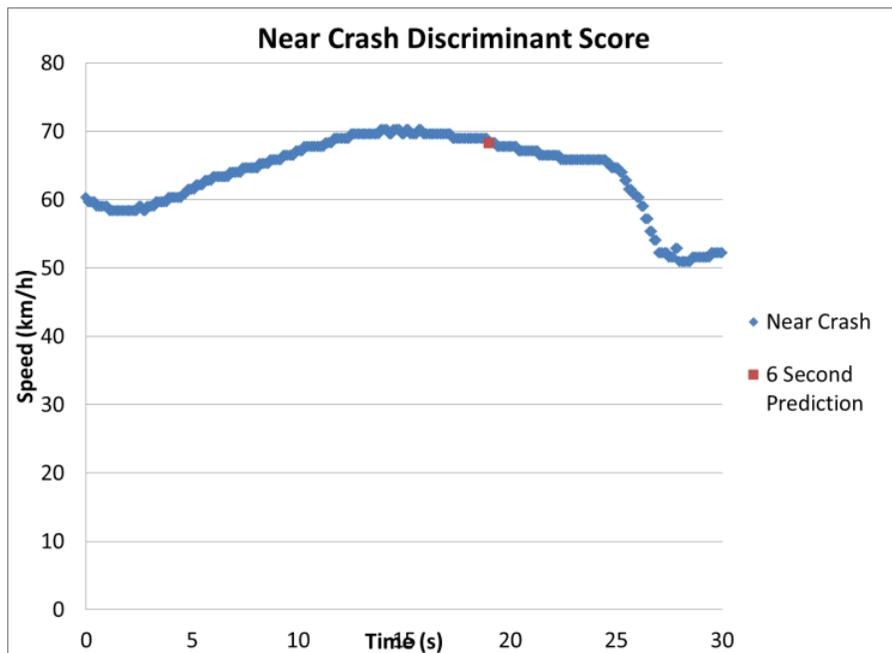


Figure 50. Graph. Trajectory of a sample near-crash event.

Table 30 shows the discriminant coefficients for Driver m. The acceleration coefficients show that the discriminant score for this driver relies on the longitudinal acceleration more than on the lateral acceleration. Driver 103-A showed more balance between the two accelerations, which

shows that every person has unique behaviors. Figure 51 shows the percentage of misclassification associated with each set of discriminant coefficients.

Table 30. Discriminant coefficients of Driver m.

Time	Longitudinal Acceleration	Lateral Acceleration	Speed	Lane Offset	Yaw Angle	Range	Range Rate
0	-15.1	-2.3	0.64	-0.02	14.9	-0.02	0.21
1	12.1	2.8	0.39	0.04	-9.0	0.01	-0.10
2	-26.8	0.0	0.32	0.02	-13.3	-0.01	0.01
3	-32.3	-0.1	0.66	-0.01	2.6	-0.02	0.05
4	-28.7	0.9	0.33	0.00	4.2	-0.01	0.02
5	-30.6	0.2	0.31	0.00	-1.8	-0.02	0.04
6	-27.4	1.0	0.21	0.00	0.7	-0.01	0.03
7	-31.2	0.3	0.29	-0.01	5.8	-0.02	0.04
8	28.2	-0.7	-0.21	0.00	4.2	0.01	-0.01
9	-11.6	2.4	0.14	0.03	-21.4	0.00	-0.07
10	12.3	3.3	0.14	0.02	-14.3	0.02	-0.10
11	3.1	2.8	0.14	0.02	-19.3	0.02	-0.06
12	-10.3	3.3	0.14	0.01	-1.5	0.00	-0.01
13	-7.1	0.0	0.14	0.01	-5.7	0.01	0.05
14	-8.5	-3.4	0.14	0.00	-7.6	0.01	0.08
15	5.1	-1.7	0.15	0.02	-28.9	0.01	-0.02

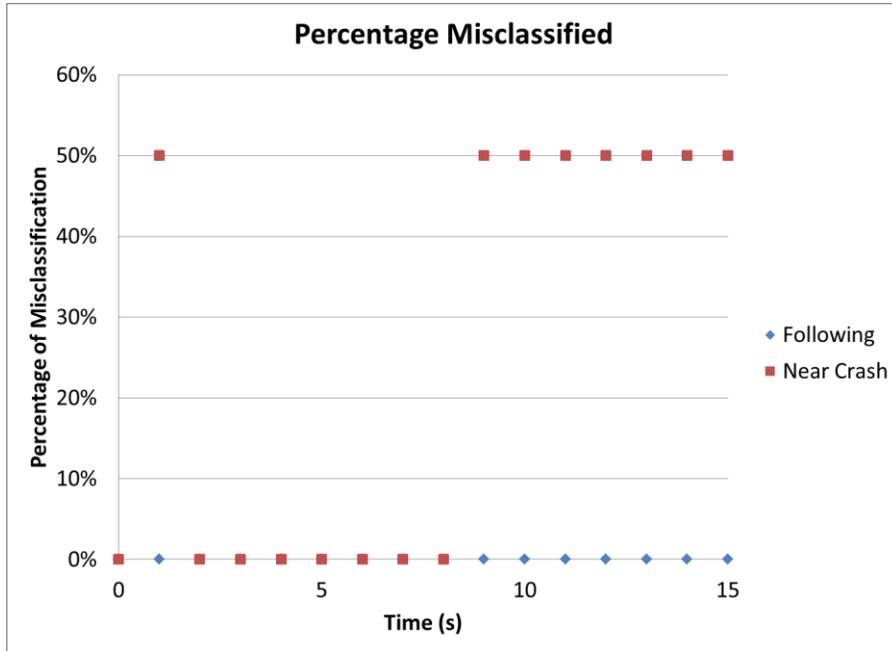


Figure 51. Graph. Percent misclassification for Driver m.

Figure 52 shows the discriminant scores for the near-crash data and normal car-following data of Driver m. These scores are calculated by multiplying the discriminant coefficients by the appropriate data values and then taking the summation. There is a separation between the normal car-following periods and the near-crash events in Driver m at most time steps until time step 9. There, the analysis begins to show a misclassification of one of the near-crash events as normal driving behavior. To avoid this error, the coefficients at time step 7 should be considered optimal.

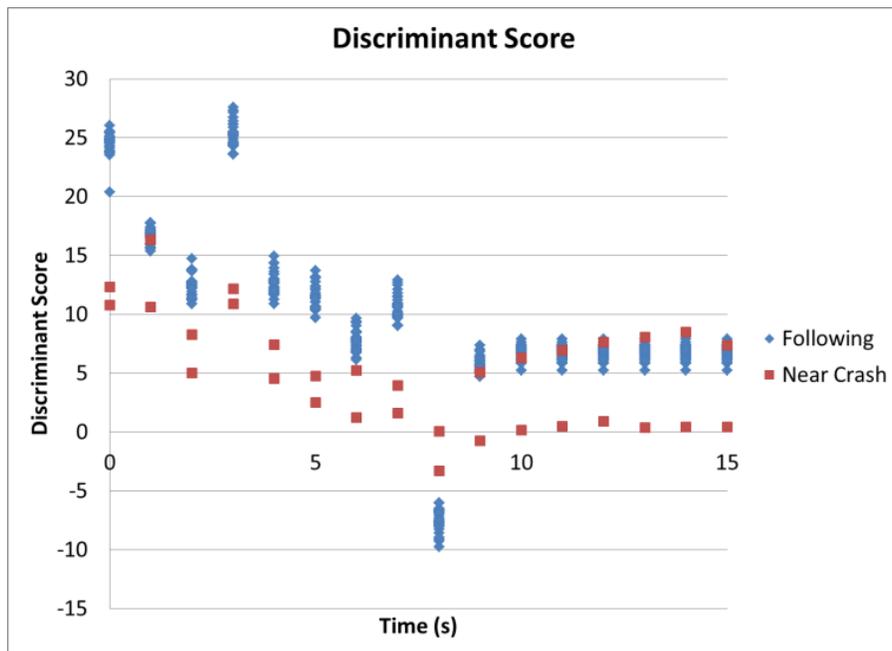


Figure 52. Graph. Discriminant scores for Driver m.

Figure 53, Figure 54, and Figure 55 show a further breakdown of the discriminant score for Driver m into three of the seven terms: longitudinal acceleration, range, and range rate. Just like Driver l, the graphs support the finding that the resulting discriminant score is dependent upon the specific combination of the input variables that the discriminant coefficients provide.

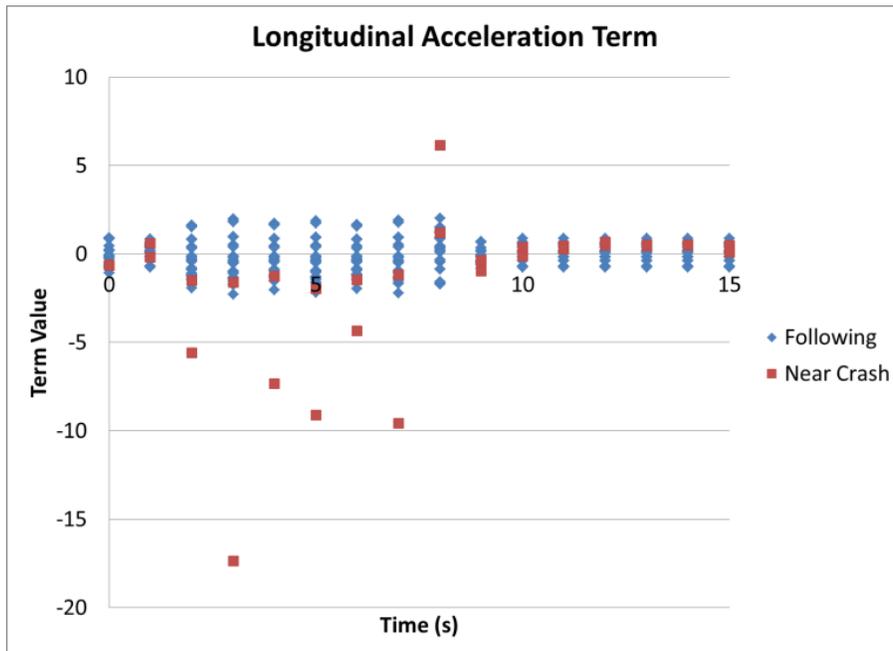


Figure 53. Graph. Longitudinal acceleration term for Driver m.

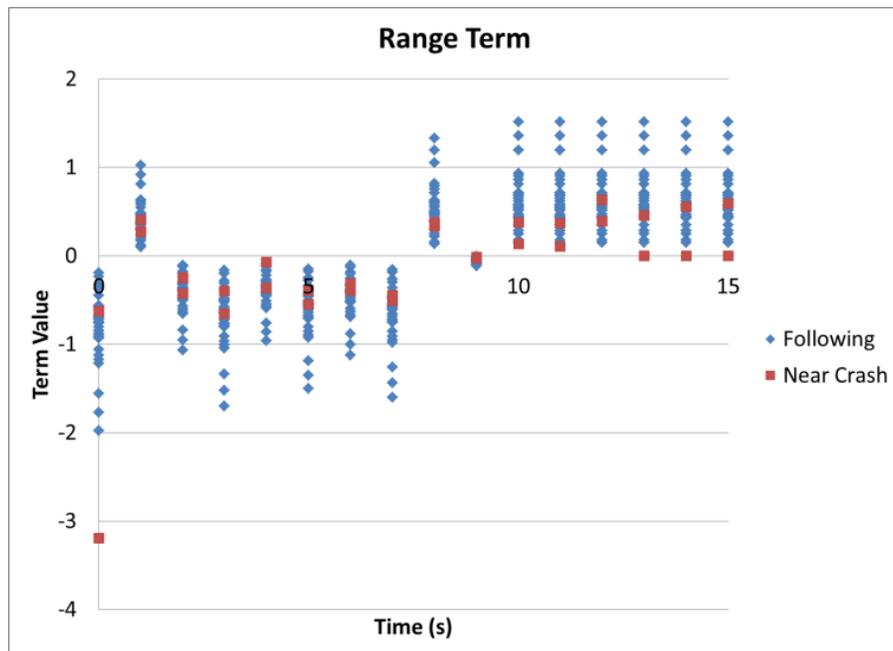


Figure 54. Graph. Range term for Driver m.

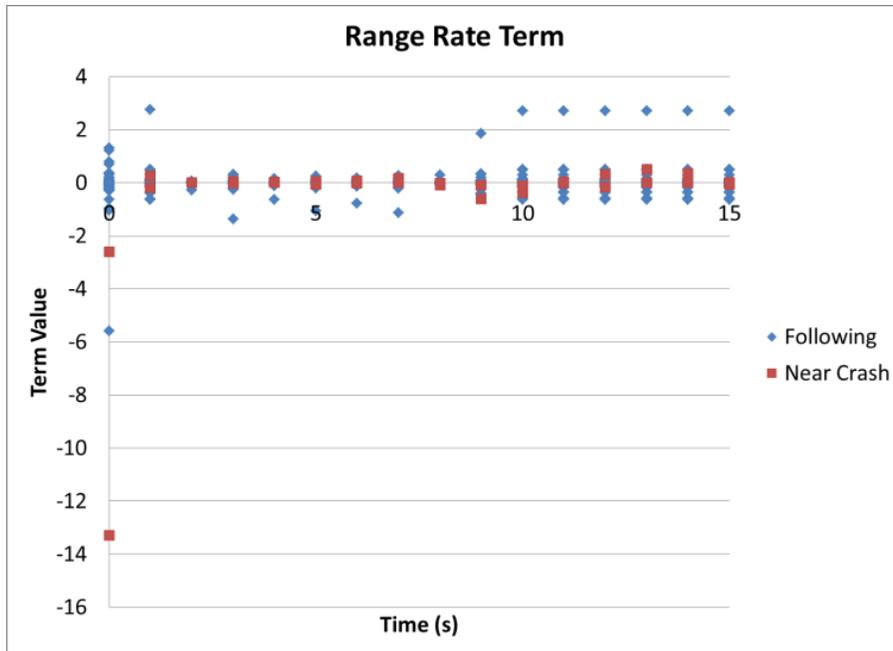


Figure 55. Graph. Range rate term for Driver m.

Figure 56 shows the resulting discriminant scores when the coefficients of the 7-s time step are applied to a full car-following period. Figure 57 shows the results of when the same coefficients are applied to the data from a near-crash event. In Figure 58 the scores remain between approximately 10 and 12, while in Figure 59 the scores decrease immediately before the event occurs. Also, the discriminant scores from 10 to 12 s in Figure 58 appear to be similar to the scores seen in the normal car-following period. This means that a deterioration of the discriminant score can be seen as the data before the near-crash event transitions from safe to unsafe. Figure 58 shows the trajectory during the near-crash event as the driver decelerates suddenly to avoid a collision. Figure 58 also shows that the 7-s prediction given by the discriminant scores appears well before the reaction of the driver.

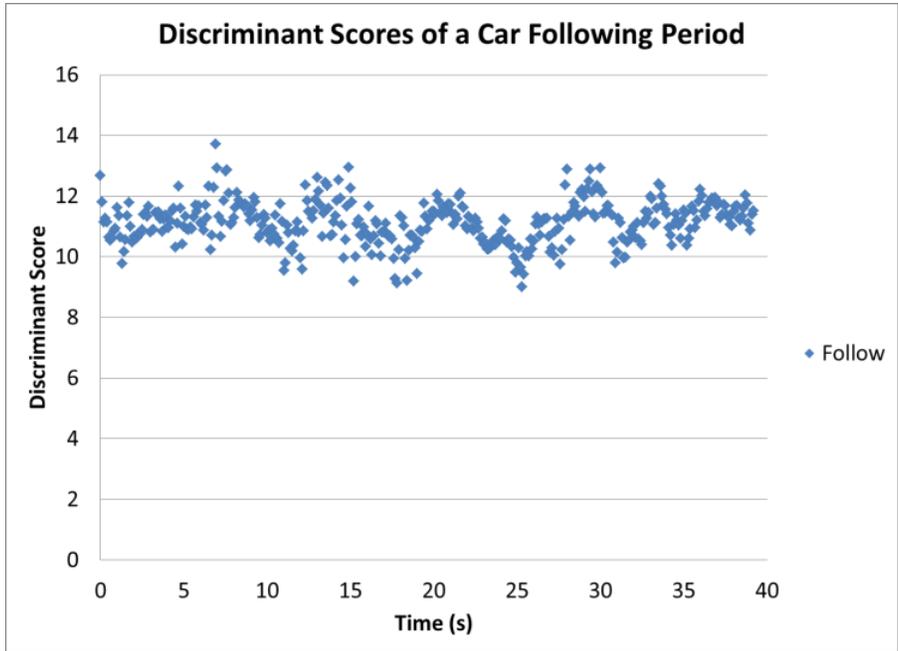


Figure 56. Graph. Discriminant scores of a sample car-following period.

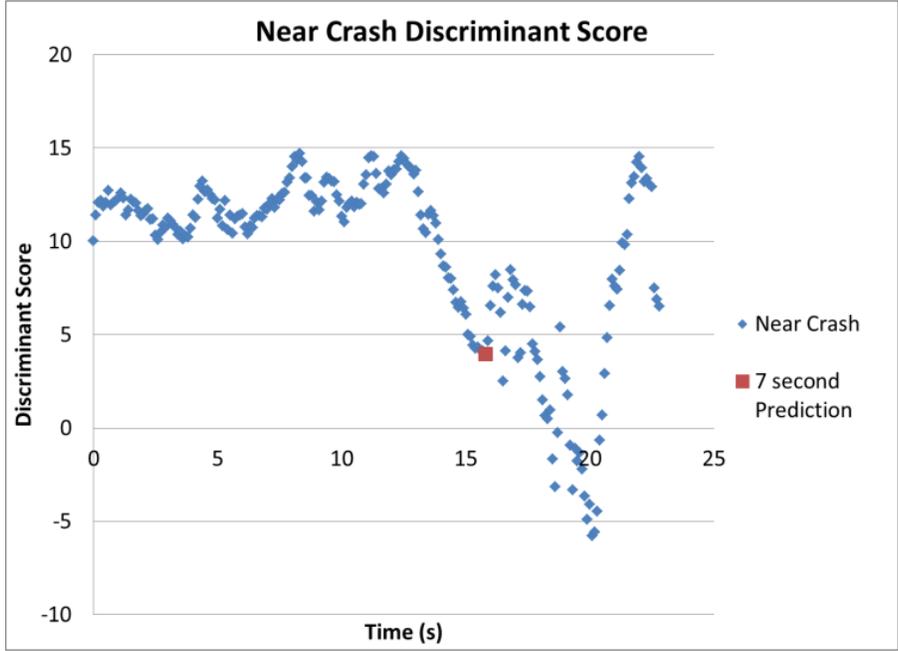


Figure 57. Graph. Discriminant scores of a sample near-crash event.

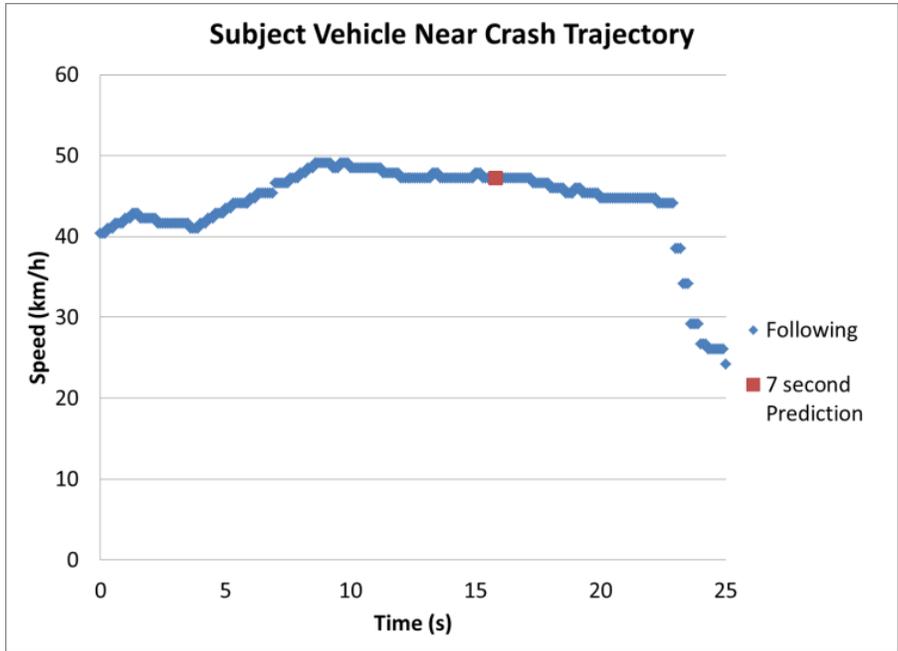


Figure 58. Graph. Trajectory of a sample near-crash event with the 7-s crash prediction point highlighted.

Figure 59 and Figure 60 show the results of attempting to use the coefficients of one driver for another driver. Both figures show that using a different set of parameters that was not optimized for that driver results in losing the clear distinction between normal car-following behavior and safety-critical behavior.

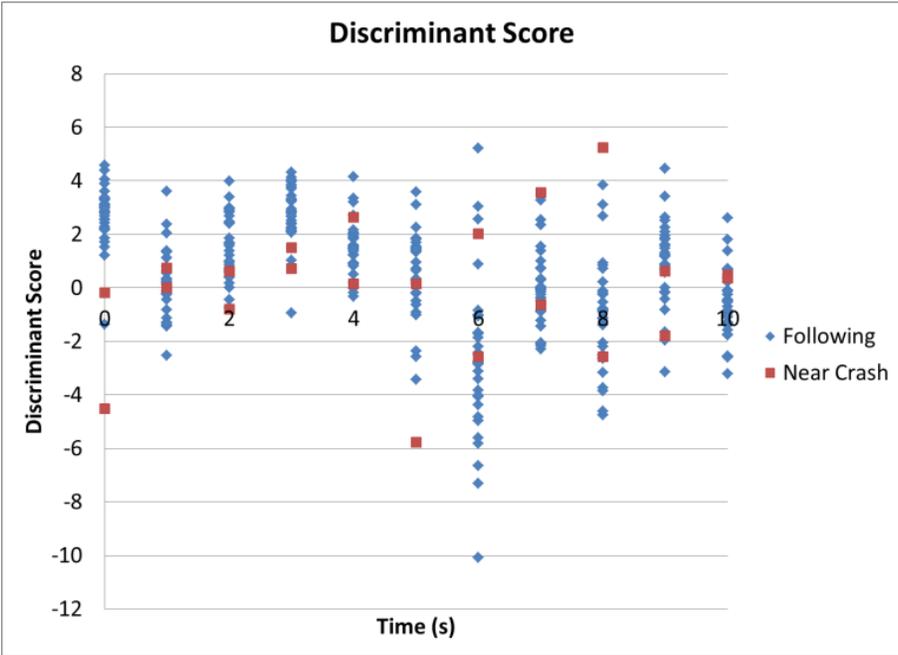


Figure 59. Graph. Discriminant score for using Driver 1's coefficients for Driver m.

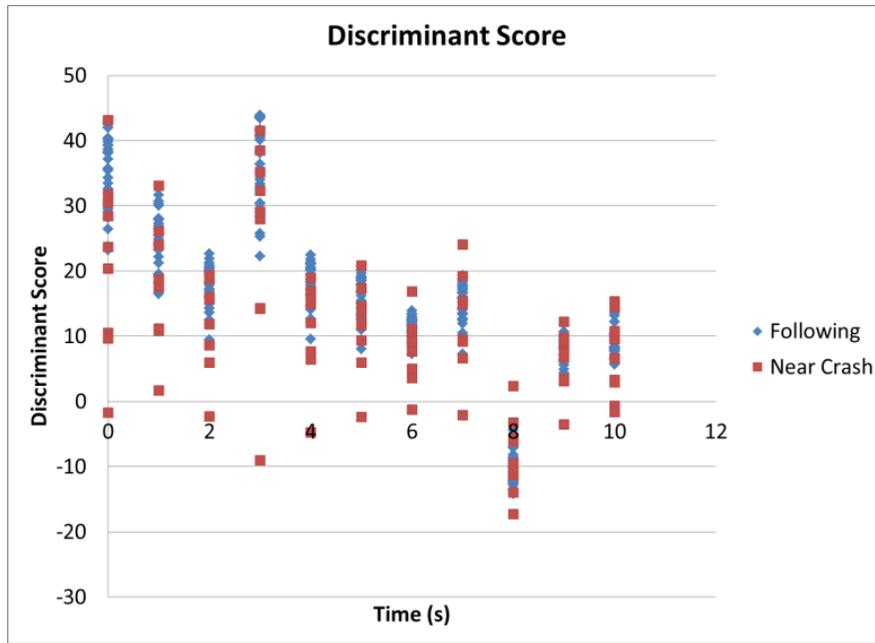


Figure 60. Graph. Discriminant scores for using Driver m’s coefficients for Driver 1.

Table 31 presents the results of applying the procedure described for the two example drivers to 8 more car drivers for a total of 10 drivers. The results show that each driver has a specific combination of the six variables and a specific threshold value that can be used to distinguish between normal driving behavior and behavior in safety-critical events. Also, the drivers’ times vary, which means that the distinction between normal driving and safety-critical behavior can be made at different times before the events for each driver. The largest time used in the analysis for some drivers was 6 s.

Table 31. Driver-specific discriminant coefficients and thresholds.

Driver	Time	Speed	Range	Range Rate	Acceleration	Yaw	Offset	Threshold
a	5.5	0.003	0.004	0.117	-13.221	74.676	-0.002	-4.525
b	6	0.013	-0.032	0.032	12.760	9.265	0.006	-7.458
c	6	0.153	-0.002	0.141	-8.460	-40.239	-0.051	0.538
d	1	0.027	0.013	0.190	-0.097	46.726	0.004	-2.728
e	6	0.087	0.005	0.243	5.184	-7.990	0.003	0.406
f	6	0.014	0.012	-0.093	-9.140	149.098	0.024	-4.317
g	6	0.056	0.003	-0.206	9.727	-73.285	0.035	-1.685
h	4.5	0.033	0.017	0.473	24.979	13.641	-0.008	0.719
i	6	0.170	-0.006	0.092	-32.342	0.000	0.227	-27.000
j	1.5	0.126	0.017	0.360	-1.852	75.317	-0.018	2.763
k	6	0.106	-0.001	-0.016	17.577	108.475	0.020	-0.504

Conclusions and Recommendations

The results support the hypothesis that using the six selected variables along with the corresponding coefficients creates a discriminant score that can accurately distinguish between safe and unsafe driving conditions. For Driver l, this distinction can be made 6 s before an event occurs.

Also, in the case of Driver m, this distinction can be made 7 s before an event occurs. The results of the discriminant analysis for Driver m arise from the specific combination of the six variables selected for analysis: longitudinal acceleration, vehicle speed, lane offset, yaw angle, range, and range rate. This means that the warning signs of an impending safety-critical event are captured and represented by these six selected variables. These warning signs are subtle changes in the conditions that go unnoticed when looking at individual variables. But using the specific combination of the six variables, as defined by the discriminant analysis, results in a noticeable warning sign.

Truck Drivers

Table 32 presents the coefficient values resulting from a discriminant analysis at different time steps before the occurrence of a safety-critical event for truck drivers. The table shows consistency in some of the coefficients over time, while the rest of the coefficients appear to be time dependent. Figure 61 shows the percentage of misclassification error at each time step. Figure 62 shows the results of applying the coefficients of the different time steps to both the near-crash and normal car-following datapoints. The results are clear: the near-crash datapoints are high negative values, and the car-following (safe) datapoints are positive values.

Table 32: Discriminant coefficients for Driver E.

Time	Longitudinal Acceleration	Lateral Acceleration	Speed	Lane Offset	Yaw Angle	Range	Range Rate
0	-6.8	13.9	0.001	0.10	-11.6	0.0003	0.12
0.5	-2.4	16.0	0.006	0.10	-13.8	0.0000	0.12
1	-0.5	15.1	0.010	0.09	-13.1	-0.0009	0.13
1.5	-3.6	10.8	0.003	0.07	40.3	-0.0016	0.13
2	0.3	9.7	0.010	0.05	36.4	-0.0028	0.14
2.5	-4.2	12.2	0.011	0.05	0.0	-0.0030	0.14
3	0.4	11.4	0.017	0.05	-4.0	-0.0038	0.14
3.5	1.2	10.6	0.018	0.04	9.3	-0.0043	0.14
4	9.6	22.0	0.008	0.11	16.1	0.0082	0.00

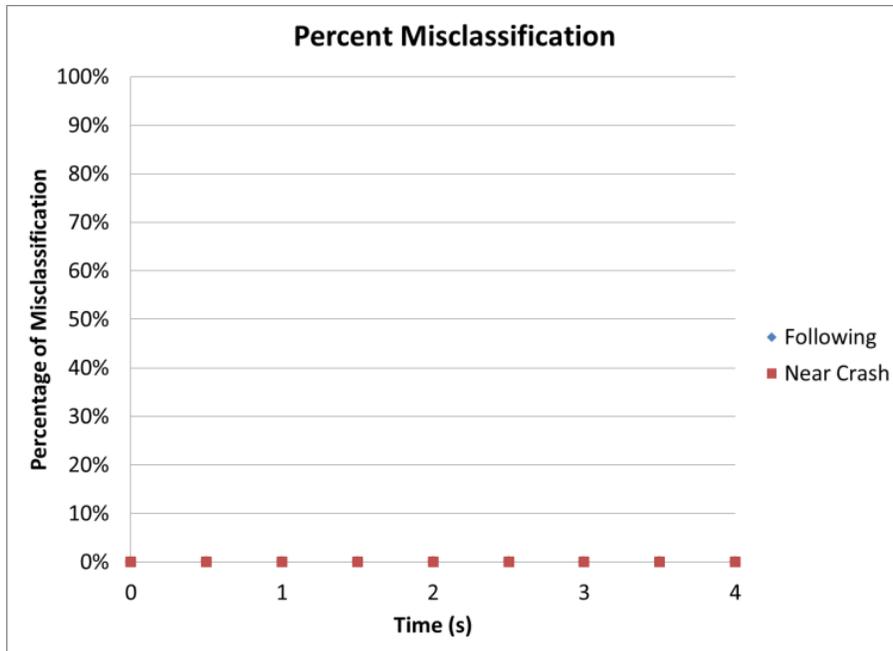


Figure 61. Graph. Percentage of misclassified datapoints for Driver E.

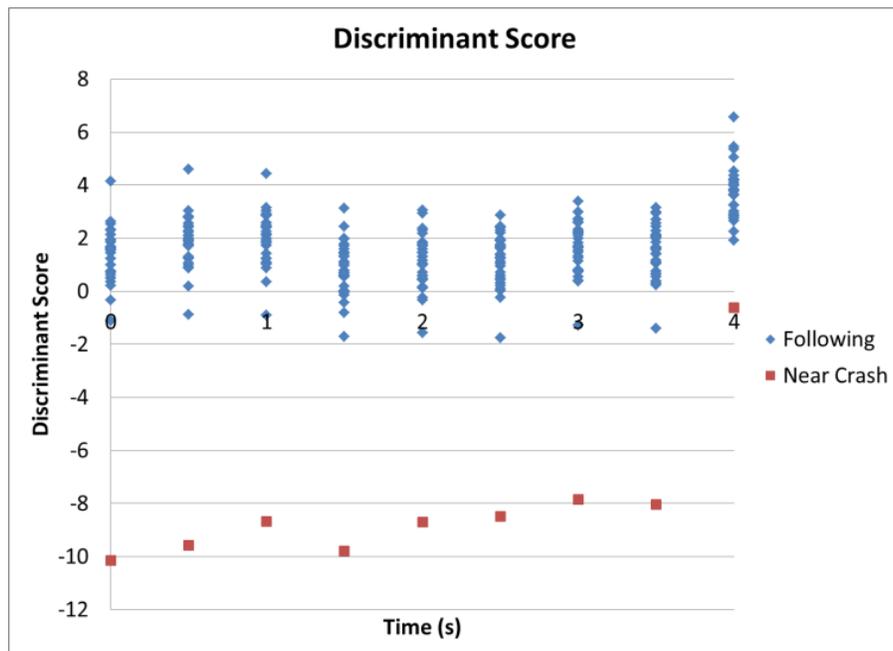


Figure 62. Graph. Discriminant scores for Driver E.

Figure 63, Figure 64, and Figure 65 show the graphs of some of the terms that constitute the discriminant score for Driver E. These graphs are the result of multiplying the input variables by the corresponding discriminant coefficient.

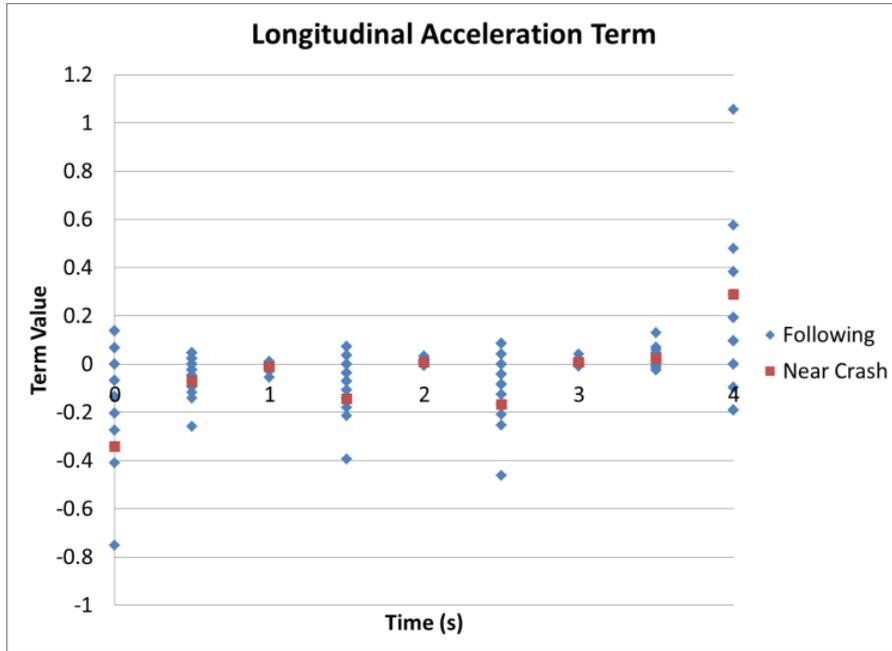


Figure 63. Graph. Longitudinal acceleration term for Driver E.

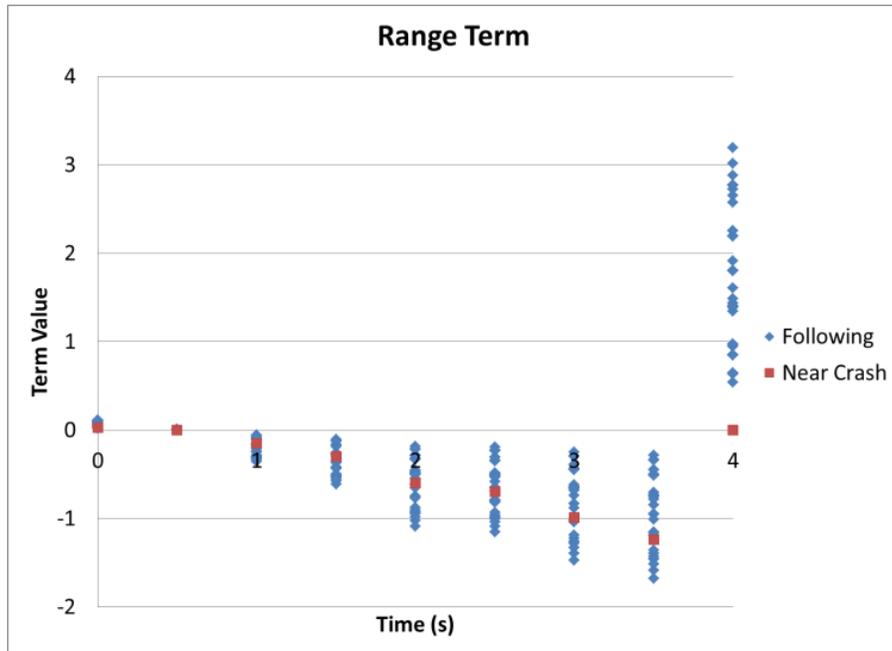


Figure 64. Graph. Range term for Driver E.

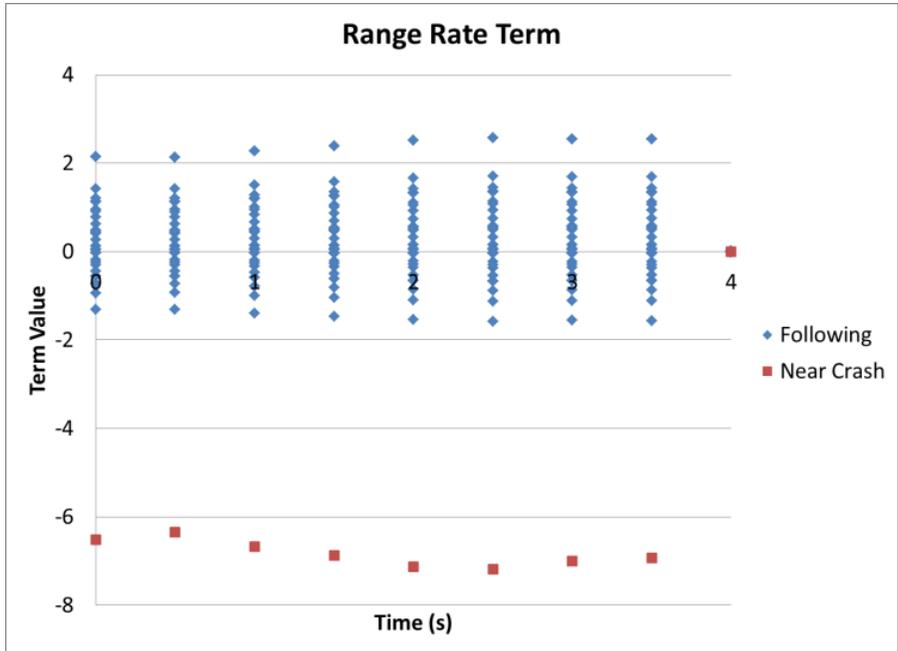


Figure 65. Graph. Range rate term for Driver E.

Table 33 shows the discriminant coefficients for Driver G. Figure 66 shows the percentage of misclassification associated with each set of discriminant coefficients.

Table 33. Discriminant coefficients of Driver G.

Time	Longitudinal Acceleration	Lateral Acceleration	Speed	Lane Offset	Yaw Angle	Range	Range Rate
0	19.7	-1.3	-0.012	0.007	-79.1	0.03	-0.05
0.5	35.4	14.8	-0.007	0.024	-52.9	0.02	0.05
1	35.7	11.5	-0.009	0.025	-59.2	0.02	0.04
1.5	37.2	9.8	-0.009	0.028	-55.2	0.01	0.05
2	33.9	4.2	-0.014	0.023	-65.9	0.02	0.01
2.5	1.1	-14.5	-0.027	-0.004	-76.8	0.04	-0.11
3	0.8	21.5	0.026	0.006	69.4	-0.03	0.13
3.5	5.2	21.3	0.026	0.011	48.6	-0.03	0.15
4	12.1	11.6	0.025	0.017	27.5	-0.03	0.14
4.5	4.1	27.3	0.023	0.010	60.4	-0.03	0.15
5	-3.2	-3.2	-0.026	-0.007	-74.2	0.04	-0.08
5.5	5.9	-15.5	-0.026	-0.002	-75.2	0.03	-0.11

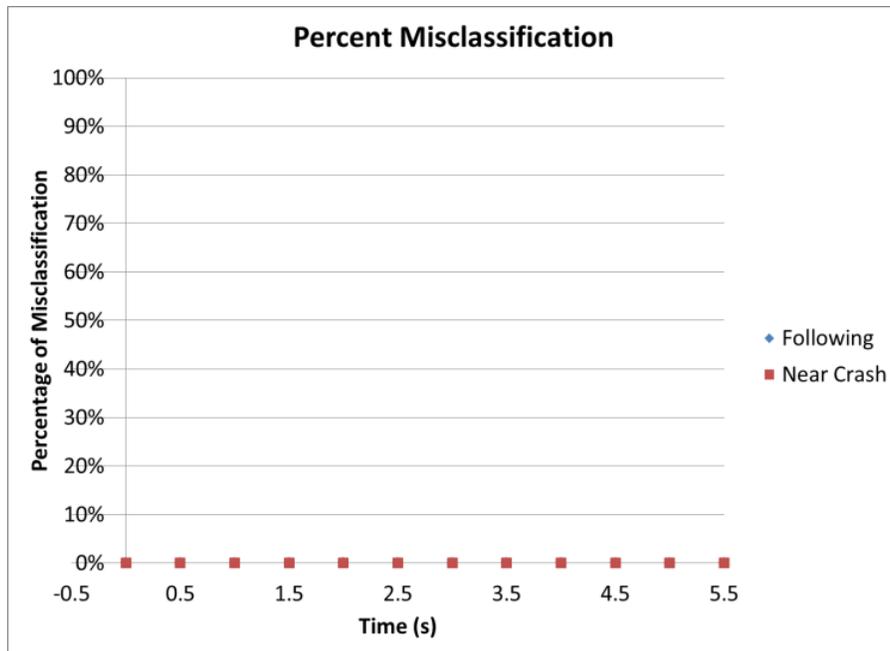


Figure 66. Graph. Percentage of misclassification for Driver G.

Figure 67 shows the discriminant scores for the datapoints from both a near-crash event and normal car-following periods. These scores are calculated by multiplying the discriminant coefficients by the appropriate data values and then taking the summation.

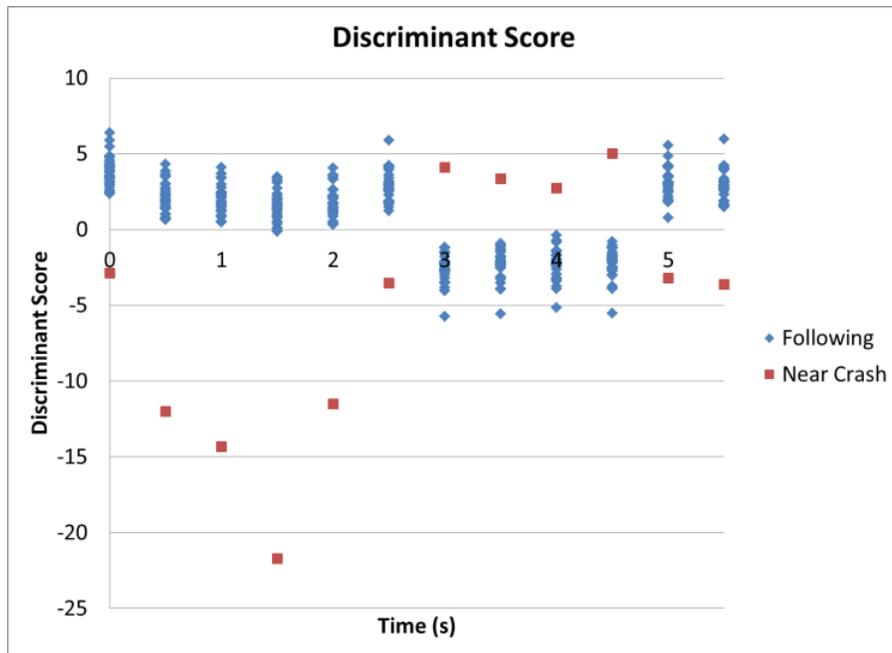


Figure 67. Graph. Discriminant scores for Driver G.

Figure 68, Figure 69, and Figure 70 show a further breakdown of the discriminant score for Driver G into three of the six terms: longitudinal acceleration, range, and range rate.

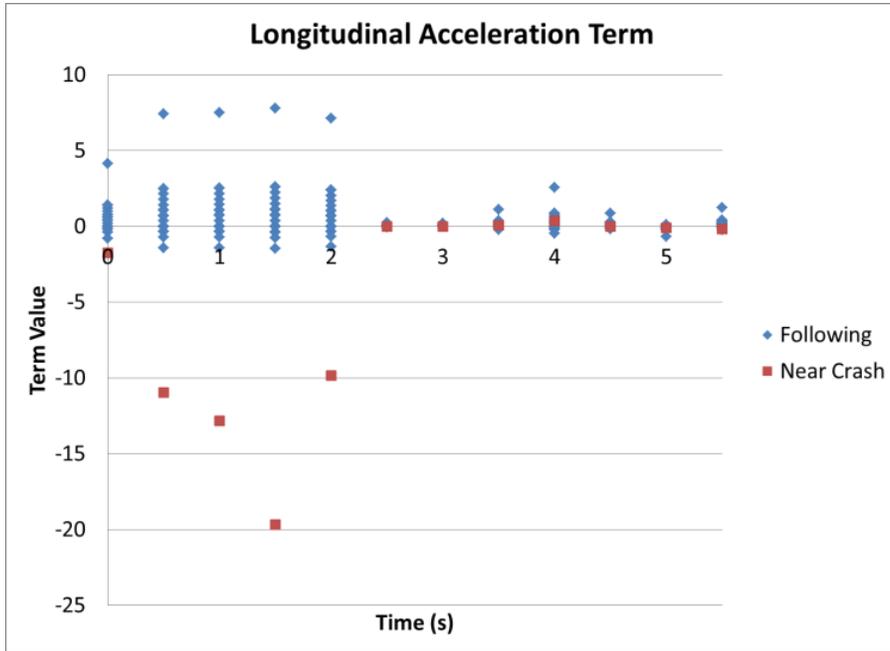


Figure 68. Graph. Longitudinal acceleration term for Driver G.

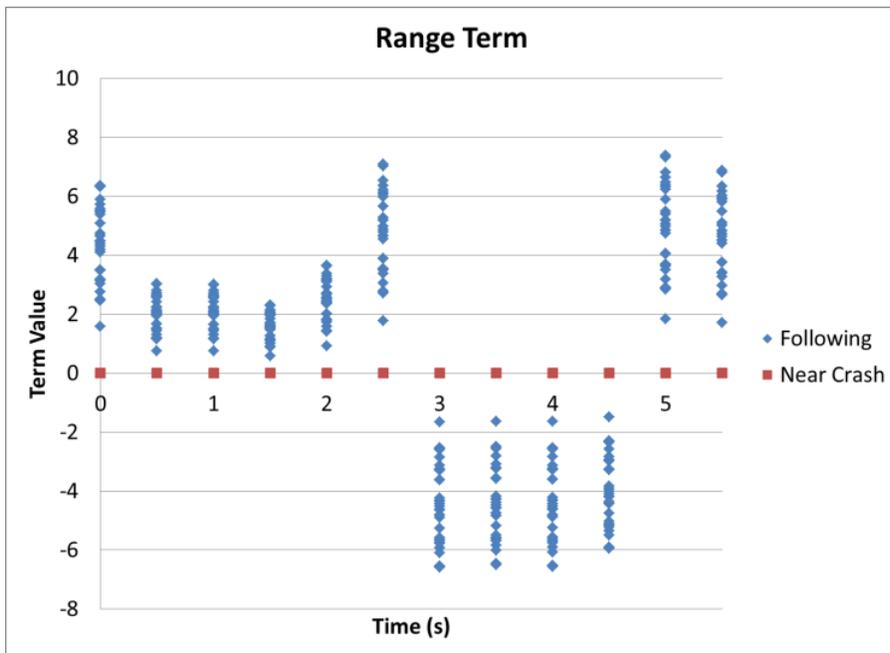


Figure 69. Graph. Range term for Driver G.

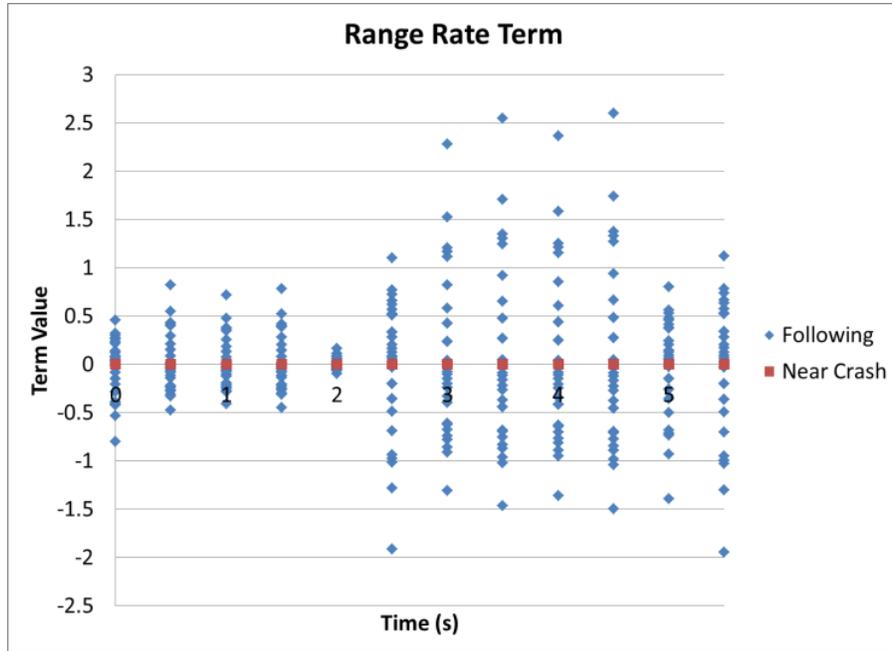


Figure 70. Graph. Range rate term for Driver G.

Table 34 shows the results of a discriminant analysis using the data from both drivers. The reason for this analysis is that truck drivers have more similarities than car drivers and may thus show common behaviors that would not exist between car drivers. The resulting coefficients from the analysis are then applied to both drivers in Figure 71 and Figure 72. The figures show that the coefficients are very accurate for both drivers until after 2 s. At that point, misclassification begins to occur for Driver G. This misclassification error is classifying a safety-critical event as normal behavior, which is not acceptable.

Table 34. Discriminant coefficients from joint analysis.

Time	Longitudinal Acceleration	Lateral Acceleration	Speed	Lane Offset	Yaw Angle	Range	Range Rate
0	19.8	9.6	0.016	0.012	-28.6	0.001	0.11
0.5	25.8	15.3	0.018	0.018	-17.4	-0.002	0.12
1	25.4	13.9	0.019	0.018	-20.7	-0.003	0.13
1.5	21.5	12.9	0.020	0.017	0.6	-0.004	0.14
2	28.4	11.8	0.019	0.019	-7.5	-0.003	0.12
2.5	0.4	6.7	-0.002	0.005	-21.0	0.002	0.13
3	2.1	0.8	0.001	0.006	-15.8	0.001	0.13
3.5	-1.2	2.8	0.001	0.003	18.8	0.000	0.13
4	-7.3	6.8	-0.014	-0.012	38.4	0.012	-0.04

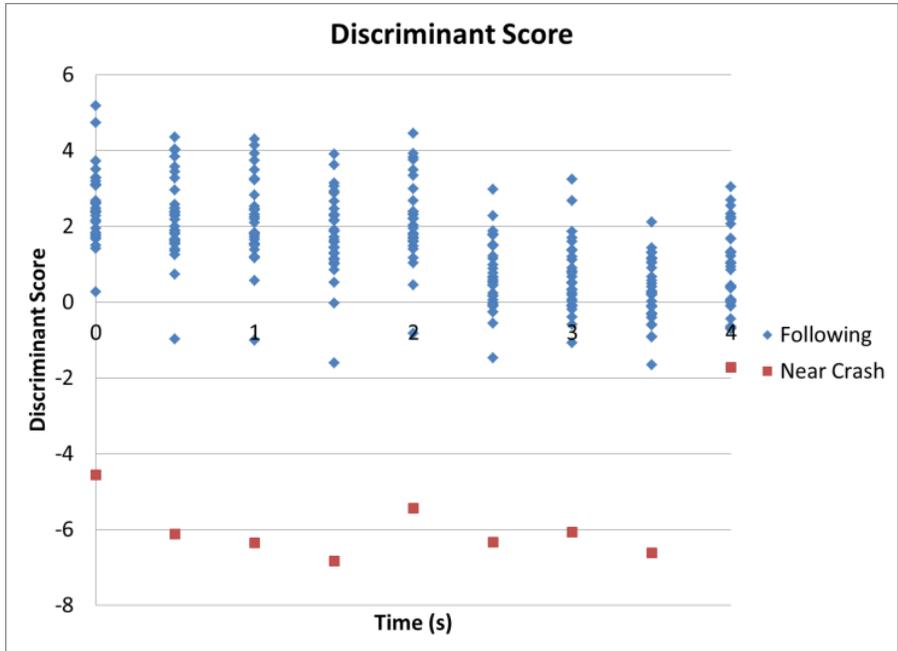


Figure 71. Graph. Discriminant score for using the joint coefficients for Driver E.

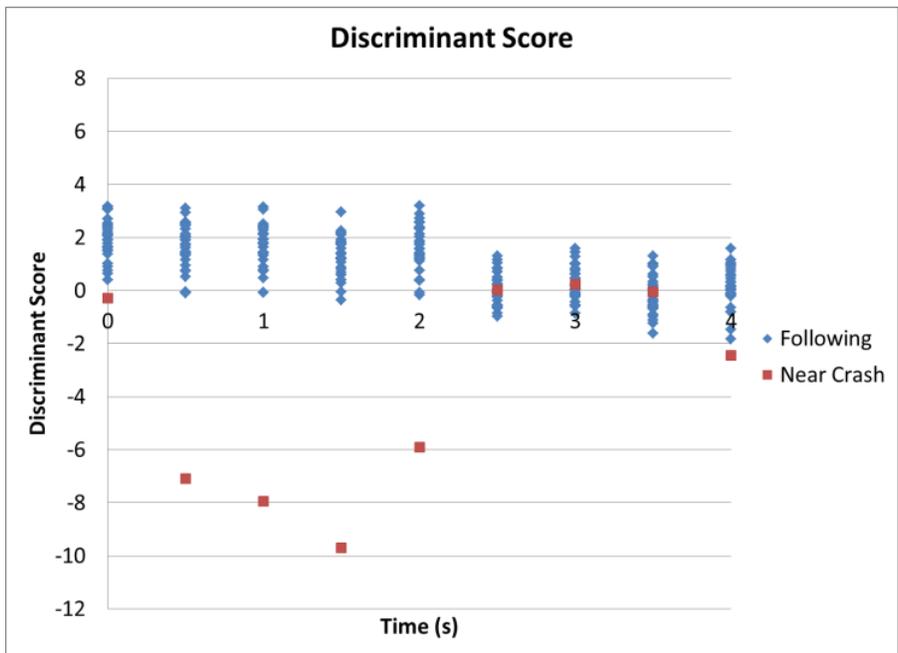


Figure 72. Graph. Discriminant scores for using the joint coefficients for Driver G.

Table 35 and Table 36 present the results of applying the described procedure for the two example drivers to eight more truck drivers from the two naturalistic truck driving studies. The results show that each driver has a specific combination of the six variables and a specific threshold value that can be used to distinguish between normal driving behavior and safety-critical event behavior. Also, the drivers' times vary, which means that the distinction between

normal driving and safety-critical behavior can be made at different times before the events for each driver. The largest time used in the analysis for some drivers was 6 s.

Table 35. 34-Truck driver-specific coefficients and thresholds.

Driver	Time	Speed	Range	Range Rate	Acceleration	Yaw	Offset	Threshold
B	6	0.584	-0.006	0.022	22.477	0.000	-0.003	-4.048
I	6	0.050	0.000	0.291	18.283	-272.421	-0.105	-15.492

Table 36. 8-Truck driver-specific coefficients and thresholds.

Driver	Time	Speed	Range	Range Rate	Acceleration	Yaw	Offset	Threshold
D	5.5	0.142	-0.005	0.186	8.525	10.129	-0.025	6.176
E	4	0.012	-0.002	0.203	-3.060	2.745	-0.051	-5.666
F	5.5	0.001	-0.013	0.123	16.541	-53.312	-0.020	-5.666
G	5.5	0.022	-0.043	0.149	-5.273	-56.211	0.004	-10.683
H	1	0.036	0.000	0.045	-11.439	-96.941	-0.050	-1.646
M	1	0.036	0.000	0.045	-11.439	-96.941	-0.050	-1.646
N	8.5	0.017	-0.012	-0.010	-21.512	-20.721	0.008	-3.299
O	1.5	0.735	-0.005	0.011	11.958	14.281	-0.049	33.157
P	3.5	0.112	-0.009	0.134	8.647	12.024	0.028	-1.956

Conclusions and Recommendations

The results show that a discriminant analysis of safety-critical event data, as compared to normal car-following data, yields a method of distinguishing between the two sets of data. This distinction uses six variables: longitudinal acceleration, vehicle speed, lane offset, yaw angle, range, and range rate. The results support the hypothesis that using the six selected variables, along with the corresponding coefficients, creates a discriminant score that can accurately distinguish between safe and unsafe driving conditions. For Driver E, this distinction can be made 4 s before an event occurs.

Also, in the case of Driver G, this distinction can be made 2 s before an event occurs. This means that the warning signs of an impending safety-critical event are captured and represented by the six variables and a specific combination of these variables as defined in the results of the discriminant analysis.

The joint discriminant analysis results show that a set of coefficients can be created that will work accurately for both drivers. This suggests that multiple drivers can be grouped together and analyzed, which simplifies situations with numerous drivers.

CHAPTER CONCLUSION

Finding the difference between normal driving behavior and safety-critical event behavior can be difficult. The statistical process of a discriminant analysis facilitates finding this difference. In

this case, a discriminant analysis is used to find the specific combination of six variables that result in a visual difference between two datasets. Also, it was found that the difference between normal behavior and safety-critical behavior extended beyond the actual event into the behavior immediately before the event. In the examples shown, there was a visible transition from safe or normal behavior to safety-critical behavior. This resulted in the ability to turn the agent behavior on and off. The trained agents offer benefits to behaviors in safety-critical events. Further, because simulation packages are designed to follow specific car-following models, the ability to turn the agent behavior on and off can be beneficial. This also means that the agents could be trained only for the safety-critical event behavior, and the car-following models can be used otherwise, which makes the application of the agents more flexible. This also bridges the gap between the current state of practice and the new trained agent method, which eases the implementation of the new method into a current simulation program, VISSIM. The next chapter provides detailed instructions of how to implement the new agent approach into VISSIM.

CHAPTER 7. AGENT IMPLEMENTATION IN VISSIM

INTRODUCTION AND PURPOSE

This chapter describes the agent implementation in VISSIM. The purpose of the demo prototype included in the deliverables of this project is to show the complete simulation chain, from generating parameter files to actually selecting the agent in VISSIM and observing its driving behavior. The following points are demonstrated:

- How to produce an agent parameter file with the APC GUI.
- How to define agent vehicles in VISSIM, including the reference to the DLL file and the parameter file.
- What settings to adjust in VISSIM to have “pure” agent behavior that is not manipulated by internal VISSIM plausibility checks.
- Behavior of the agent model in case of an event that is close to an event observed in reality.

OVERALL WORKFLOW

Figure 73 shows the overall flow of the implementation framework. A parameter file is used to get the model parameters into VISSIM using the driver model DLL. The parameters used are those generated by the techniques described in the previous chapters. The implementation framework utilizes one parameter file for each agent; in other words, for 20 agents there will be parameter files *ParameterFile_Agent_1.APC*, *ParameterFile_Agent_2.APC*, ... *ParameterFile_Agent_20.APC*.

In VISSIM, vehicle types can be defined that use a specific DriverModel.DLL as an external driver model. (See the description in the next section and the VISSIM manual section 5.3.1. Note that the parameter file referred to in the manual and in the VISSIM Vehicle Type dialog does not conform to the parameter files used in this project.)

The DriverModel.DLL writes a specific output file with information about the driver state, and so on. To allow the files to be written properly, there must be a specific DLL for each agent in the model. The name of the DLL is part of the output file name (which will be saved in the current working directory).

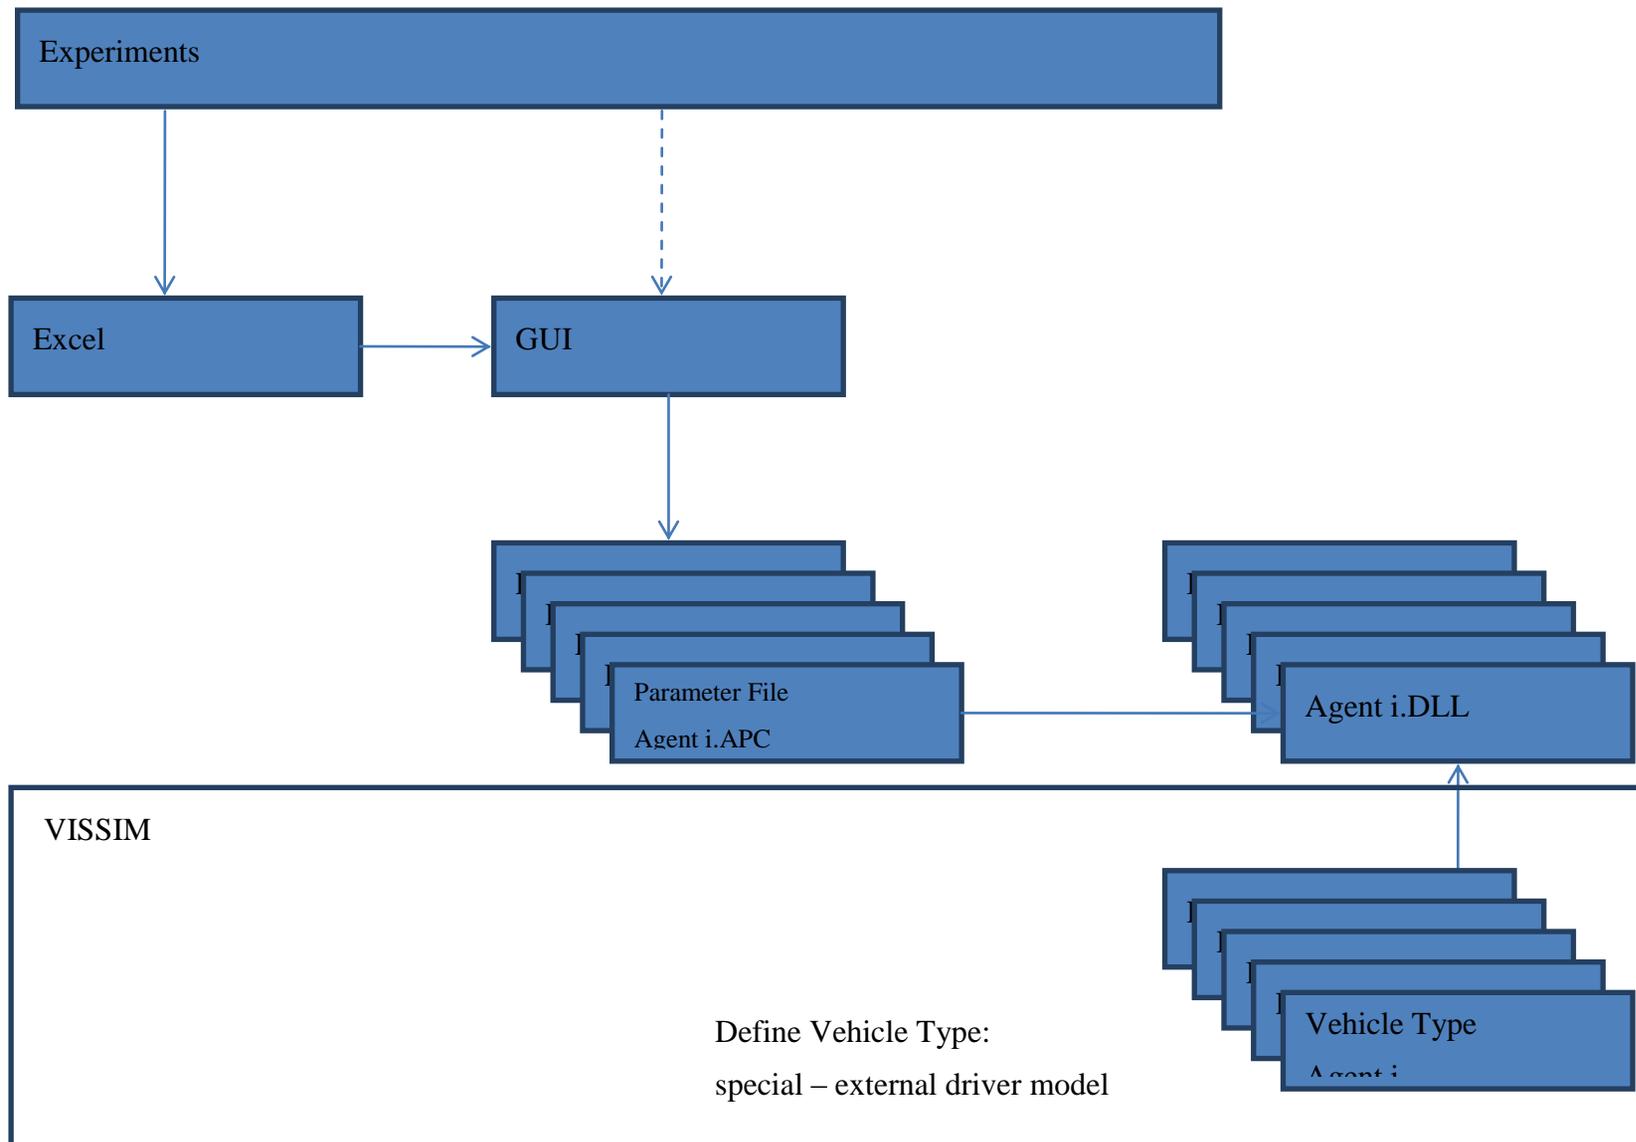


Figure 73. Chart. Overall workflow.

VISSIM SETTINGS

Vehicle Types

Two new vehicle types are defined: agent and leader. The respective driver model DLL and parameter files should be selected as shown in Figure 74 (file locations have to be adjusted to the local machine).

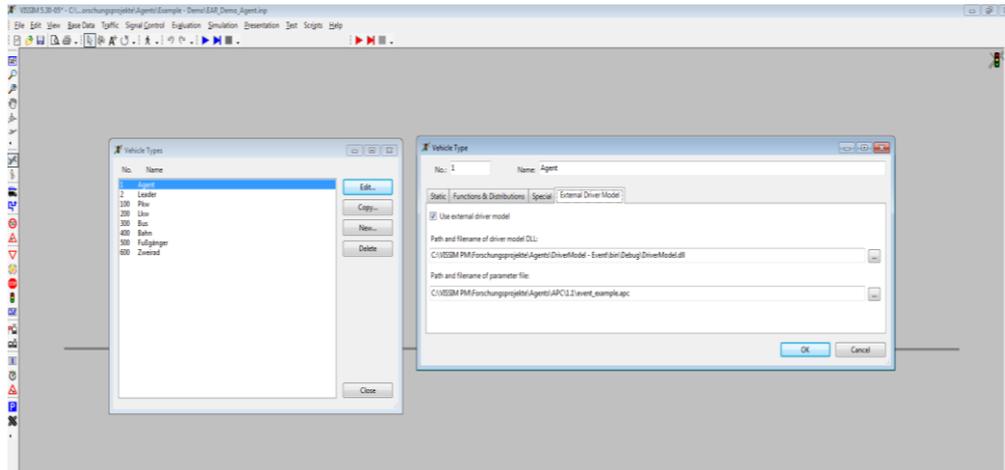


Figure 74. Screen shot. Specification of vehicle types to use the agent model.

2D/3D Model Distribution

For the exact interpretation of the lane angle set by the Agent Model DLL, the geometry of the 2D model used for the agent vehicle type needs to be adjusted. To make sure everything works properly, perform the following steps:

1. Check the Model Distribution that is used for the Agent vehicle type (e.g., 10 PKW) (Figure 75).

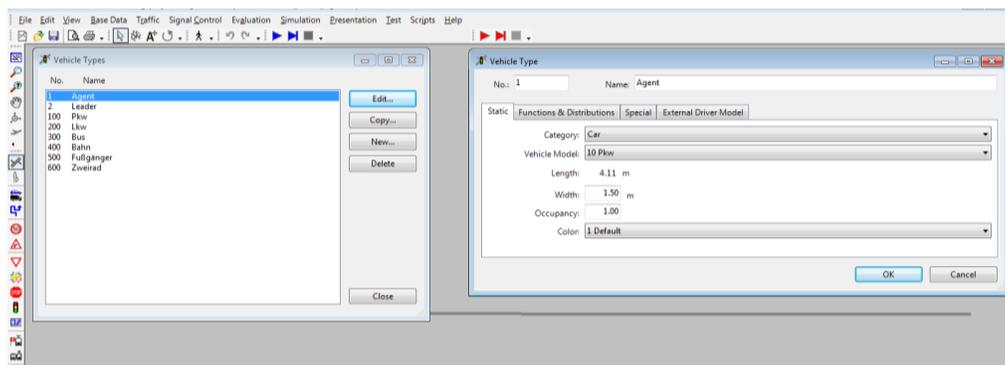


Figure 75. Screen shot. Model distribution for agent vehicle type.

2. Open the Model Distribution dialog and edit the Model Distribution (Figure 76).
 - a. Use only one model in the distribution (i.e., set the share for all models except one to zero).

- b. For that model, move the rear axle to the rear end of the vehicle (i.e., set Rear Axle and Rear Joint to the same value as the vehicle length, e.g., 4.110 m).

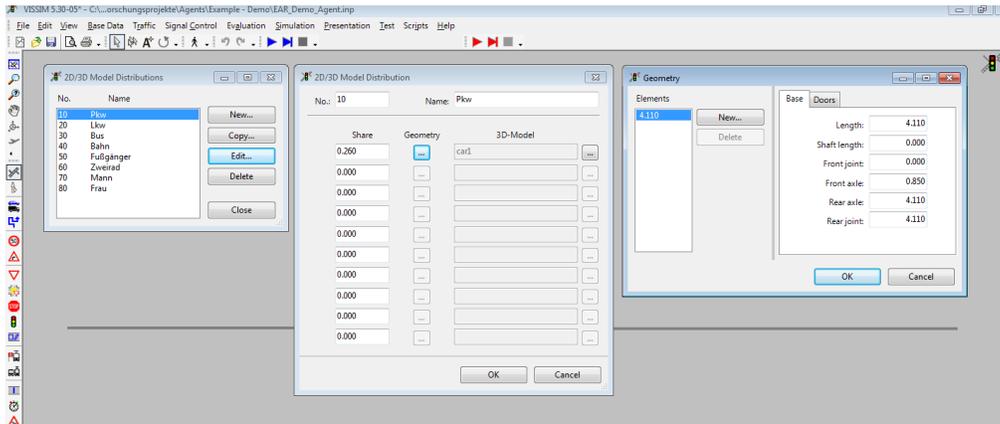


Figure 76. Screen shot. Model distribution dialog.

DRIVING BEHAVIOR

In order for the lane angle to be exchanged between VISSIM and the Agent Model DLL, the “Observe vehicles on next lane(s)” option in the Driving Behavior Parameter Sets dialog (Figure 77) needs to be activated (even if there is only one lane).

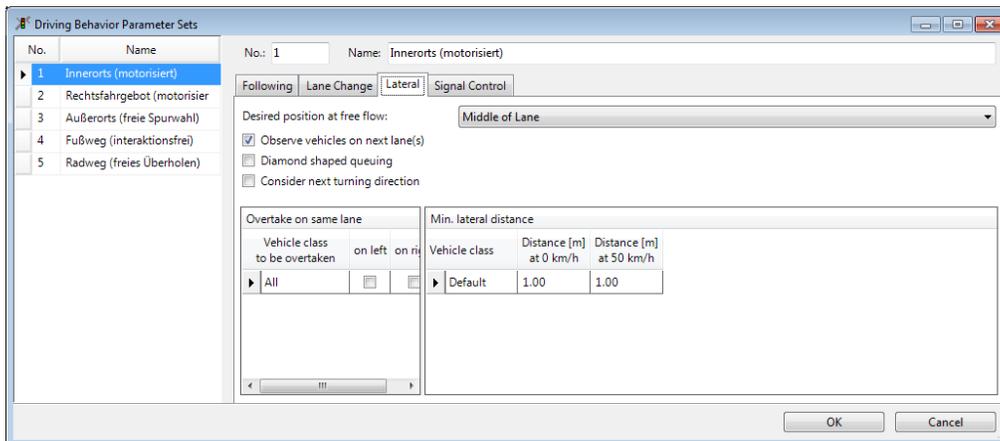


Figure 77. Screen shot. Lateral driving behavior.

ACCELERATION AND DECELERATION FUNCTIONS

In VISSIM, acceleration and maximum deceleration functions for the vehicles controlled by the DLLs have to be introduced. This is because VISSIM checks if the acceleration and deceleration proposed by the DLL lies within the limits defined in VISSIM by these functions. Namely, VISSIM sets the applied acceleration to the minimum and maximum desired acceleration and the applied deceleration to the minimum and maximum desired deceleration. The maximum acceleration and deceleration functions are named in the “DLL.” The maximum acceleration function is defined as $+7 \text{ m/s}^2$ at any speed, and the maximum deceleration function is defined as

-100 m/s² at any speed. These limits are clearly out of the feasible range and should not be used for normal simulations. In the demo, they are used to produce accident-like decelerations (leader) and to show the pure agent behavior for the follower.

PARAMETER FILE *.APC

Each driver is represented by an .apc file. This section briefly explains the GUI and the parameter file content and format.

GUI

There are two ways to get the data into the parameter file:

1. The model parameters are entered in an Excel file (with the same structure of the file *NFACRL structure thresholds and weights_example-March 2011.XLSX*). As the model is implemented in the spreadsheet, the response of the model (outputs Acceleration and Lane Angle) can be checked using different state variables. In this way, the Excel file can act as kind of a plausibility check for the model parameters. The parameters can be imported to the *Agent Parameter GUI*. From the GUI, the parameter file can be written.
2. The second option is to enter the parameters directly into the GUI and export them into the parameter file. This is not recommended for tutorial purposes because it is more cumbersome and there is no way to check if the model will produce meaningful output (of course, it can be checked once the parameter file is loaded to the driver model DLL and the parameters are used in VISSIM).

In addition to the model parameters, there are threshold values to be defined. The thresholds define when the driving behavior switches from standard VISSIM (i.e., Wiedemann model) to agent-based behavior and vice versa.

To parameterize the agents, a GUI is provided (as illustrated in Figure 78 and Figure 79). The GUI gives access to all user-definable parameters of the agent, as well as to a global threshold value that determines in what situations the agent model should be used instead of the standard VISSIM driving behavior.

To parameterize the agent, the GUI provides tabs for the different model weights and thresholds, as well as a text field for the global threshold.

There is a button to read the values directly from an Excel file with the same structure of the file *NFACRL structure thresholds and weights _example-March 2011.XLSX*.

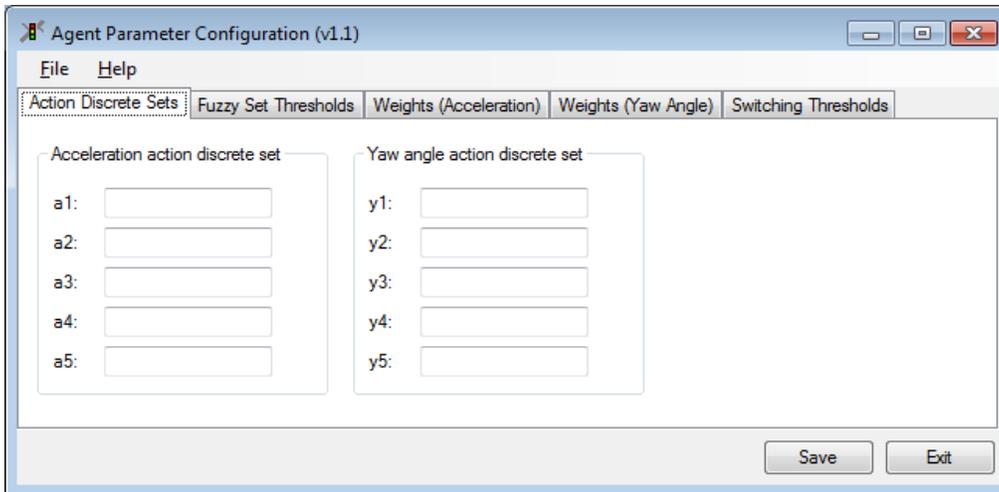


Figure 78. Screen shot. Agent parameter GUI—action discrete sets.

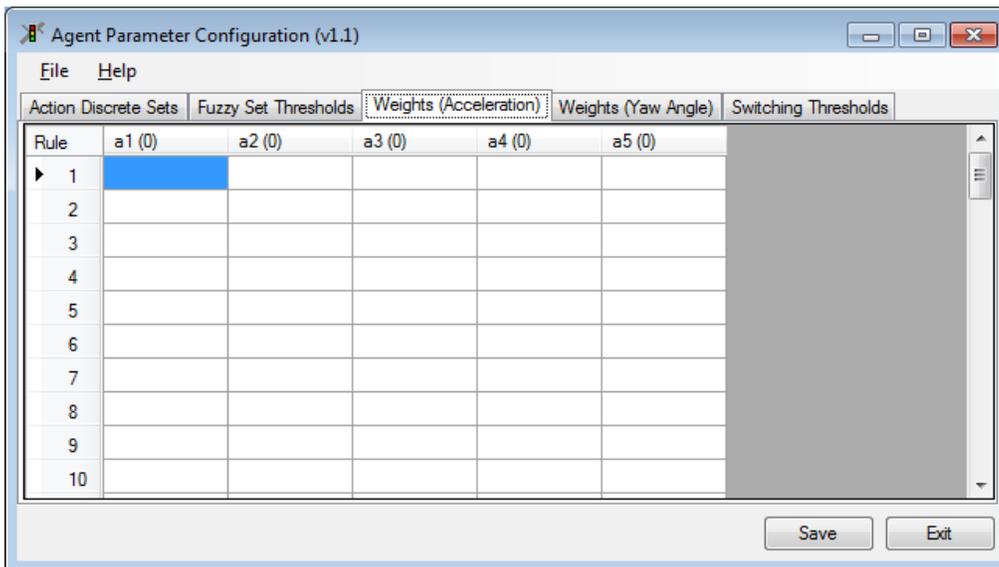


Figure 79. Screen shot. Agent parameter GUI—acceleration weights.

For this project, there is one type of agent model—in other words, the number of neurons, the rules, the layers, and the calculation logic are set in the DLL. Any changes in the model structure (such as adding a new layer) will require an updated driver Model DLL and a changed GUI to set the parameters.

Parameter Contents and Format

Table 37 shows a description of the parameter file format. The parameter file is a comma-separated text file, and all numbers are double precision format. There should be one parameter file per agent (i.e., for 20 agents there will be files *parameterFile_[n].APC*, with $n = 1 \dots 20$).

Table 37. Parameter file description.

Description	Values (in blocks, separated by commas); see APC documentation
Fuzzy Set Layer Thresholds	Fuzzy Set Layer Thresholds
thresholdSpeed:	lowerthresholdSpeed, upperthresholdSpeed
thresholdRange:	lowerthresholdRange, upperthresholdRange
thresholdRangeRate:	lowerthresholdRangeRate, upperthresholdRangeRate
thresholdAcceleration:	lowerthresholdAcceleration, upperthresholdAcceleration
thresholdLaneAngle:	lower thresholdLaneAngle, upper thresholdLaneAngle
thresholdOffset:	lower thresholdLaneAngle, upper thresholdLaneAngle
AccelerationDiscreteSet:	a1,a2,a3,a4,a5
LaneAngleDiscreteSet:	la2,la2,la3,la4,la5
Weights W1(i,j):	W1(1,1), ..., W1(1,5), W1(2,1), ..., W1(i,j),...,W1(32,5)
Weights W2(1,j):	W2(1,1), ..., W2(1,5), W2(2,1), ..., W2(i,j),...,W2(32,5)
Switch Conditions	There will be one (or more) threshold value that determines if the agent model is switched on or off. At the moment, this is Time to Collision. It can, of course, be changed to other values.

SCENARIO DESCRIPTION "EVENT"

Description

The example presented here is based on an observed real-world scenario. It should show the behavior of the VISSIM implementation of the agent in comparison to the observed data and the behavior of the Excel implementation of the agent. The following vehicle is controlled by an agent according to the parameters provided in the example Excel file. Because external driver model *DriverMode.DLL* has to be selected, the parameter file is *Event_Example.APC*.

It should be noted that directly comparing VISSIM output to naturalistic data output can be problematic for two reasons:

1. All naturalistic data variables are observed variables, whereas the same variables are a mix of observed and calculated state variables in VISSIM. In other words, the range and range rate are VISSIM observed state variables; the others are calculated state variables. Substituting the range and range rate into the VISSIM model will not always produce the corresponding calculated variables (in VISSIM, knowing the acceleration of the follower and the speed of the leader will result in exact calculation of the range and range rate for the next time step; in naturalistic data, the leader might perform any action that will result in a different range and range rate). This might lead to discontinuity in the state-space surface. For example, if the data showed that the vehicle stopped, and the model predicts that there is acceleration > 0 , the next modeled acceleration will have to either use the observed data as the previous state or accumulate a shift in the modeled state from the observed state. In the provided Excel sheet, the model uses a mixture of observation and previous model output as state variables

for the next time step. In the VISSIM model, range and range rate are naturally simulation outputs, and they respond to the model output (i.e., if an acceleration > 0 is calculated, the vehicle accelerates regardless of any observation that is saved in the spreadsheet).

2. Another issue that should be noted is that the example results start with an initial state in terms of speed, range, range rate, acceleration, lane angle, and offset. It is not possible to create exactly that initial condition in VISSIM such that, starting from there, the results of the Excel implementation and VISSIM can be compared. That is because of limitations of the DLL interface and the discrete time steps. In other words, it is possible to generate a state that is very close to the initial state described in the observed data, but it will not be the very same initial state.

These two facts (the former more than the latter) suggest that the Excel model and the VISSIM model should be compared by other means:

1. Generate a test case in VISSIM (close to the observed event) and run the simulation with the agent model switched on if the follower is as close to the leader as in the observed event.
2. Use the state variables from VISSIM (instead of reality) to calculate the model's output in Excel.
3. If the two outputs are identical (or close to identical, due to numerical reasons), it shows that the VISSIM implementation behaves like the Excel implementation.

How to Run

Run COM Script

For the leading vehicle to perform the observed trajectory, the COM interface is used to set the speed of the leading vehicle. To run the simulation, the button "Run VISSIM Example" (see Figure 80 below) in the sheet "VISSIM-COM" must be clicked.

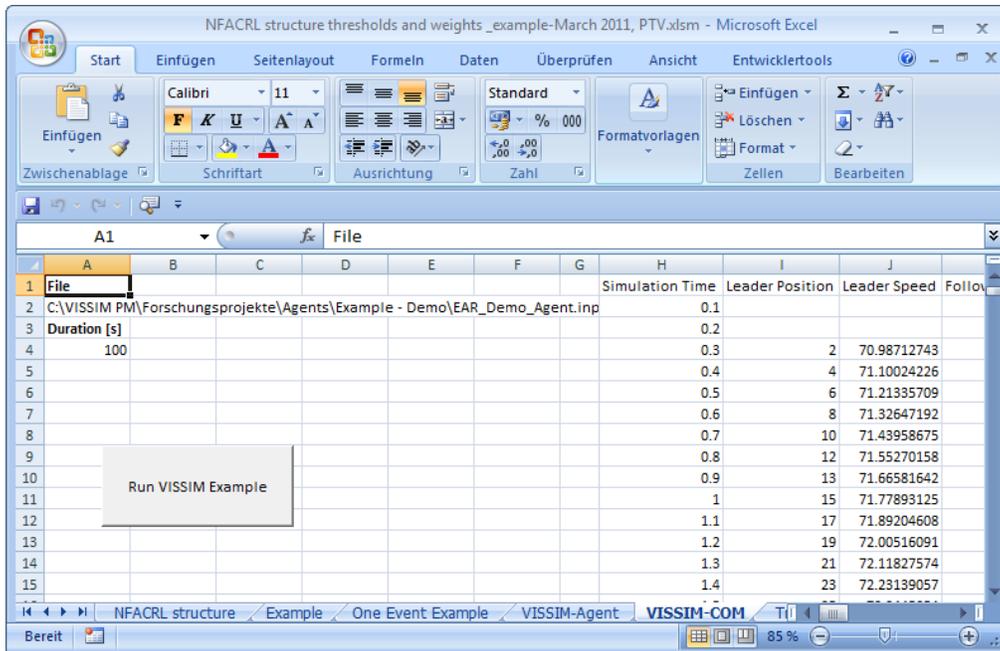


Figure 80. Screen shot. “Run VISSIM example” command.

The leading vehicle drives at relatively low speed; the following vehicle enters the network 10 s later with a higher speed (approximately constant speed is assumed). The DLL and the COM script constantly check to see if the (steadily decreasing) range between the vehicles falls below a threshold (120.9 ft) that is derived from the observation data provided in the Excel spreadsheet. If this condition is fulfilled, the leading vehicle’s speed is set according to the observed values (gathered from the Excel sheet by the COM script) and the action of the following vehicle calculated based on the Agent Model DLL.

Copy Output Data from Text File to Excel

The simulation output is written to a text file “Data_Test_[Name of the used parameter File].txt” and saved in the working directory. The content of this file needs to be copied to the Excel file (mark the whole content in the text file and paste in the Excel file, sheet “VISSIM-Agent,” at cell A1).

Run the Excel Model

To run the Excel model to compare the output derived from Excel with the model output in VISSIM, click the button “Run Excel Agent based on VISSIM data” in sheet “VISSIM-Agent” as shown in Figure 81.

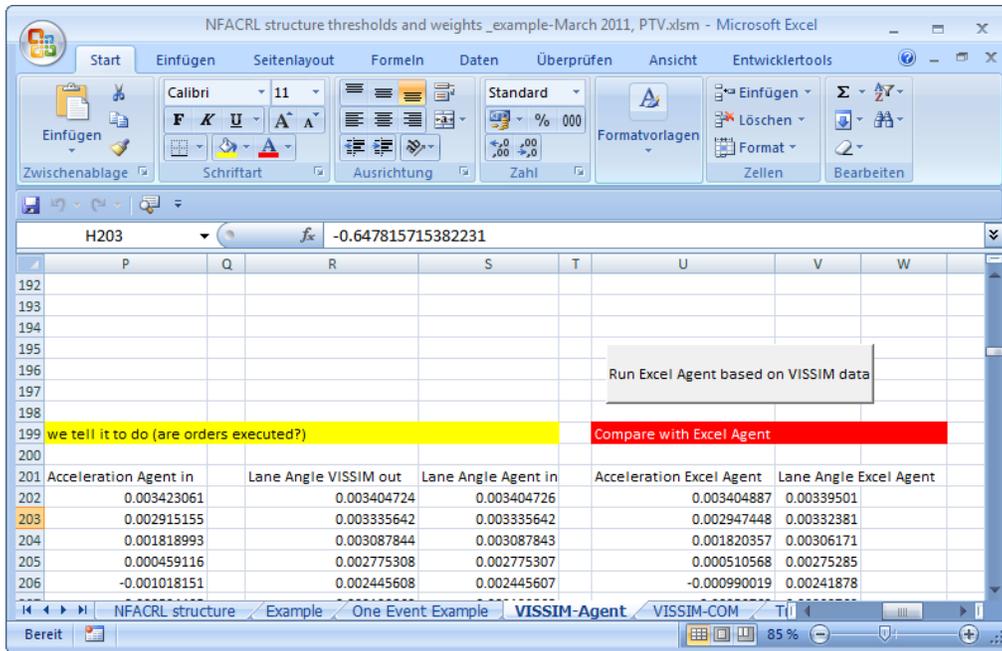


Figure 81. Screen shot. “Run Excel Agent based on VISSIM data” command.

Figure 82 through Figure 86 shown in the next section are generated automatically.

Output

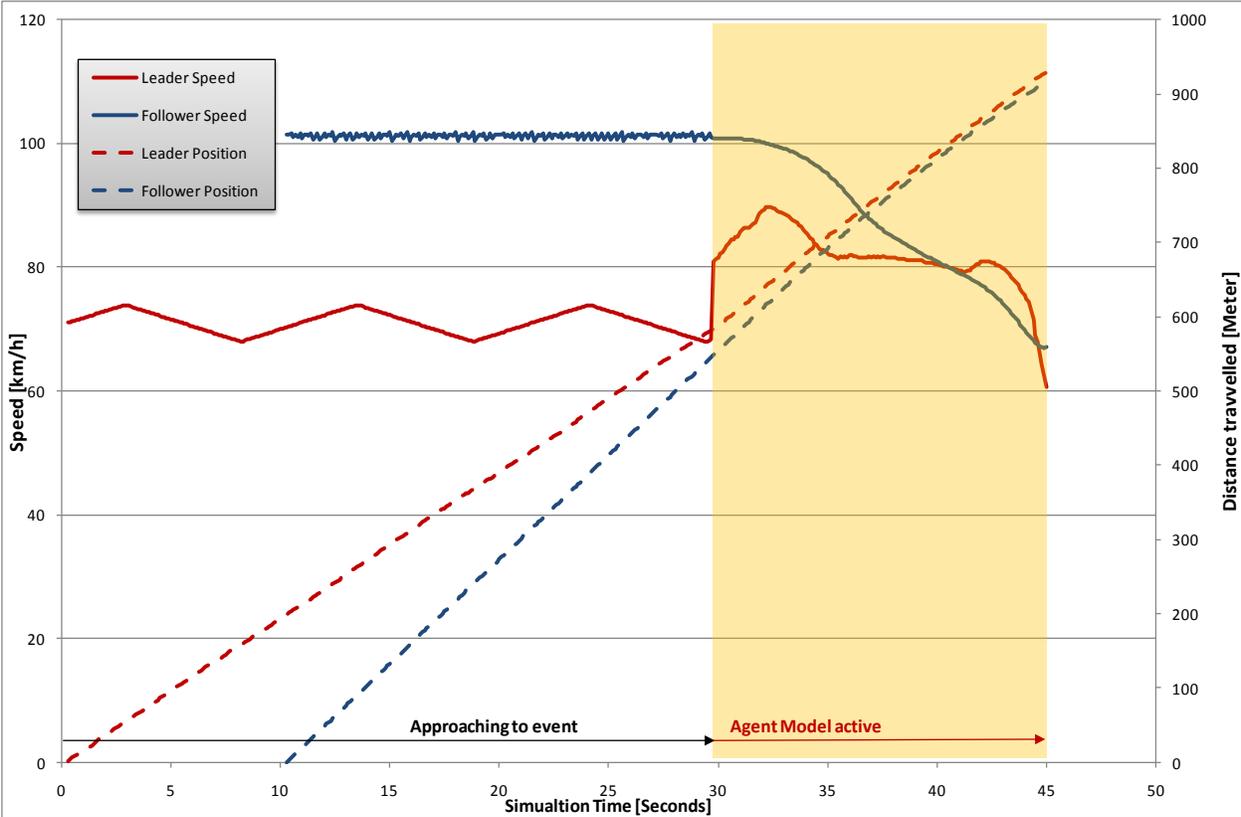


Figure 82. Graph. Comparison of speeds and trajectories of leading vehicle and agent over time.

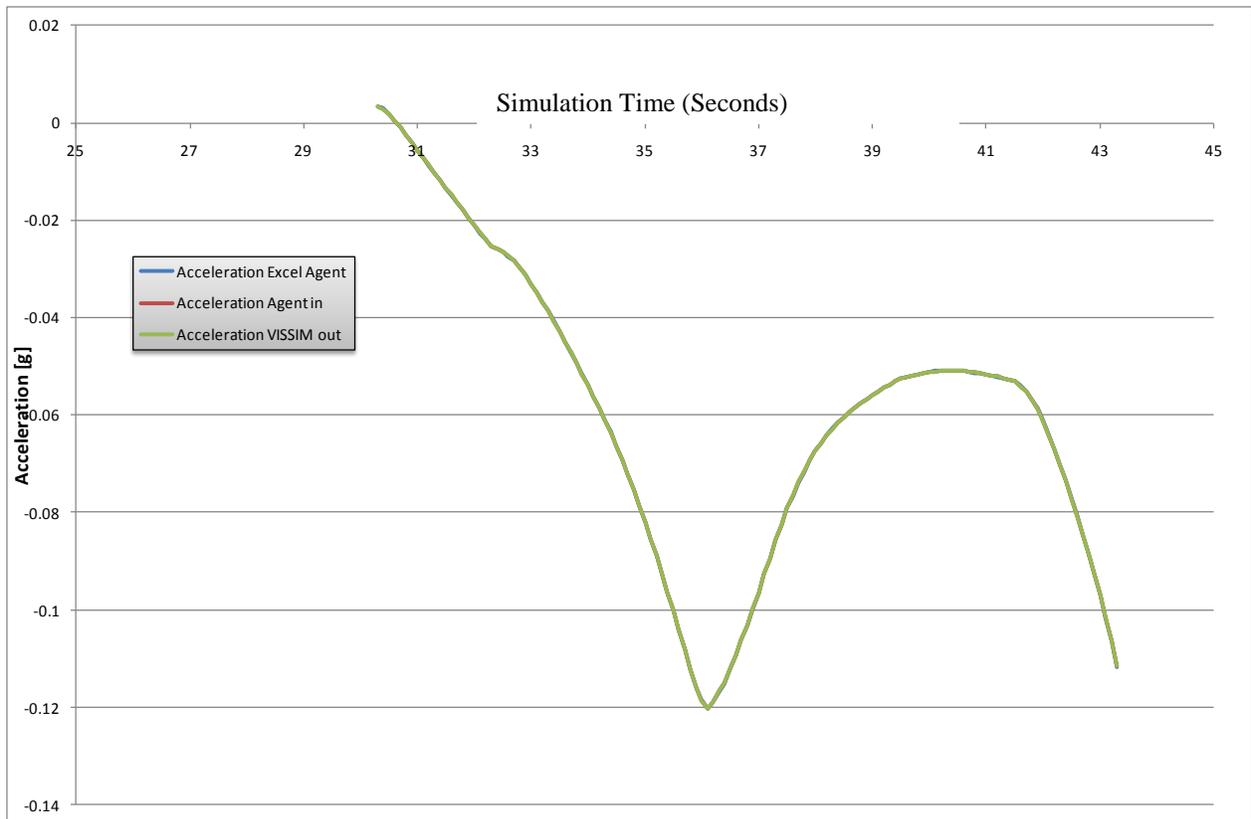


Figure 83. Graph. Comparison of acceleration over time (acceleration calculated in Excel as reference, acceleration calculated in VISSIM and performed in VISSIM to be compared).

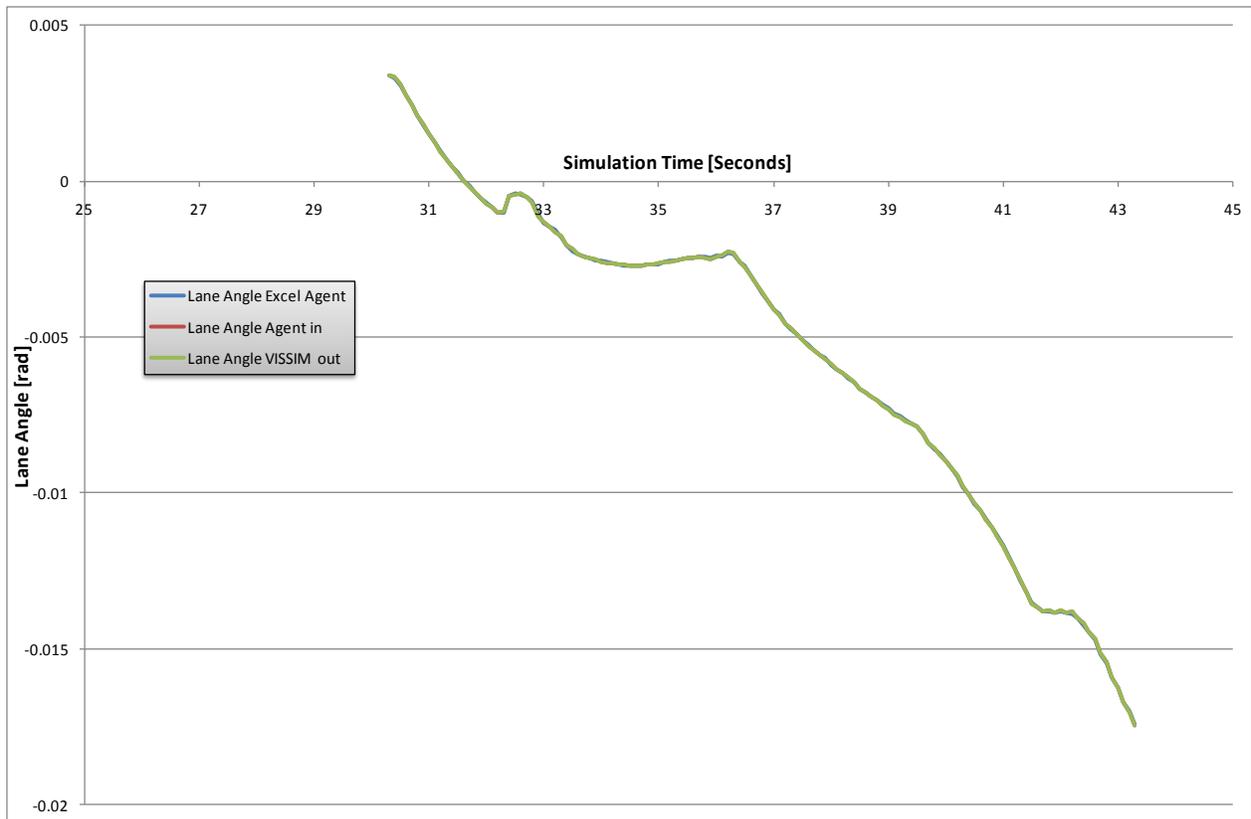


Figure 84. Graph. Comparison of lane angle over time (lane angle calculated in Excel as reference, lane angle calculated in VISSIM and performed in VISSIM to be compared).

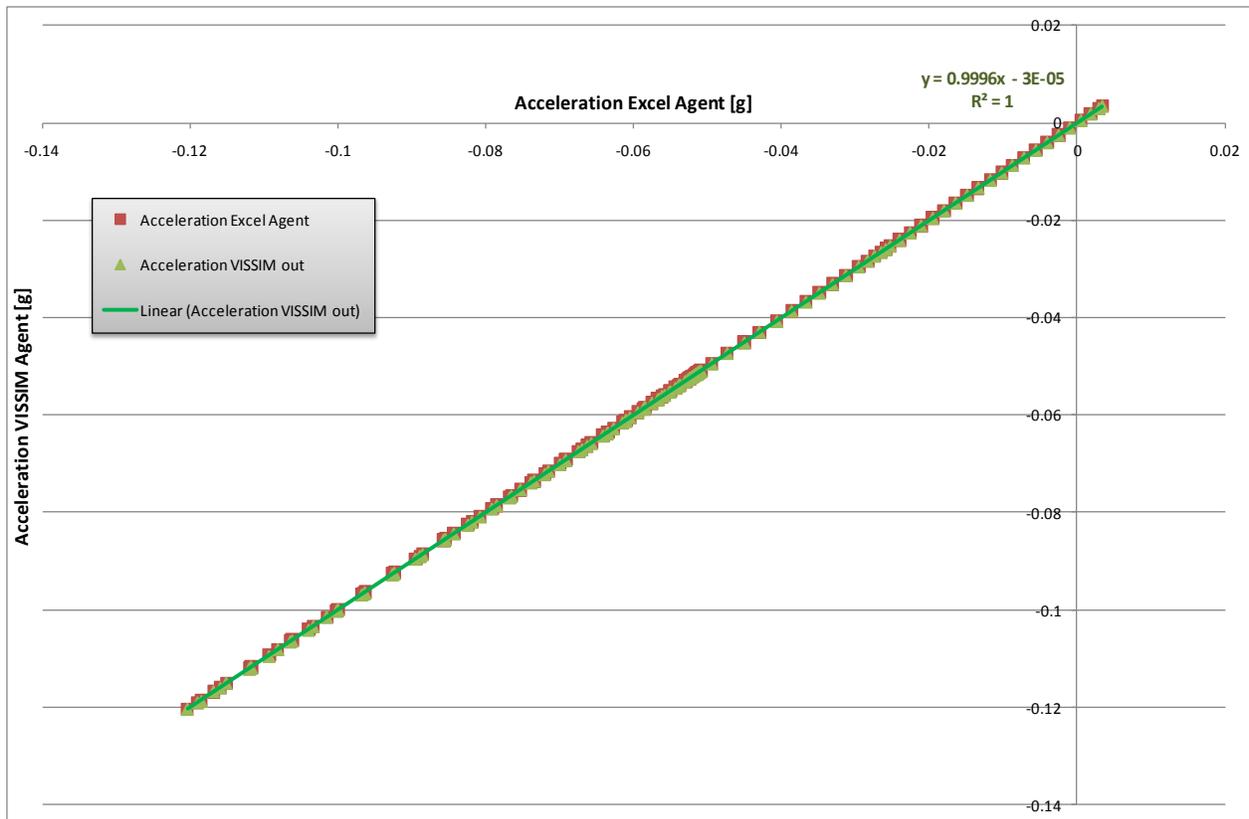


Figure 85. Graph. Acceleration scatter plot with linear regression analysis, Excel agent versus VISSIM agent.

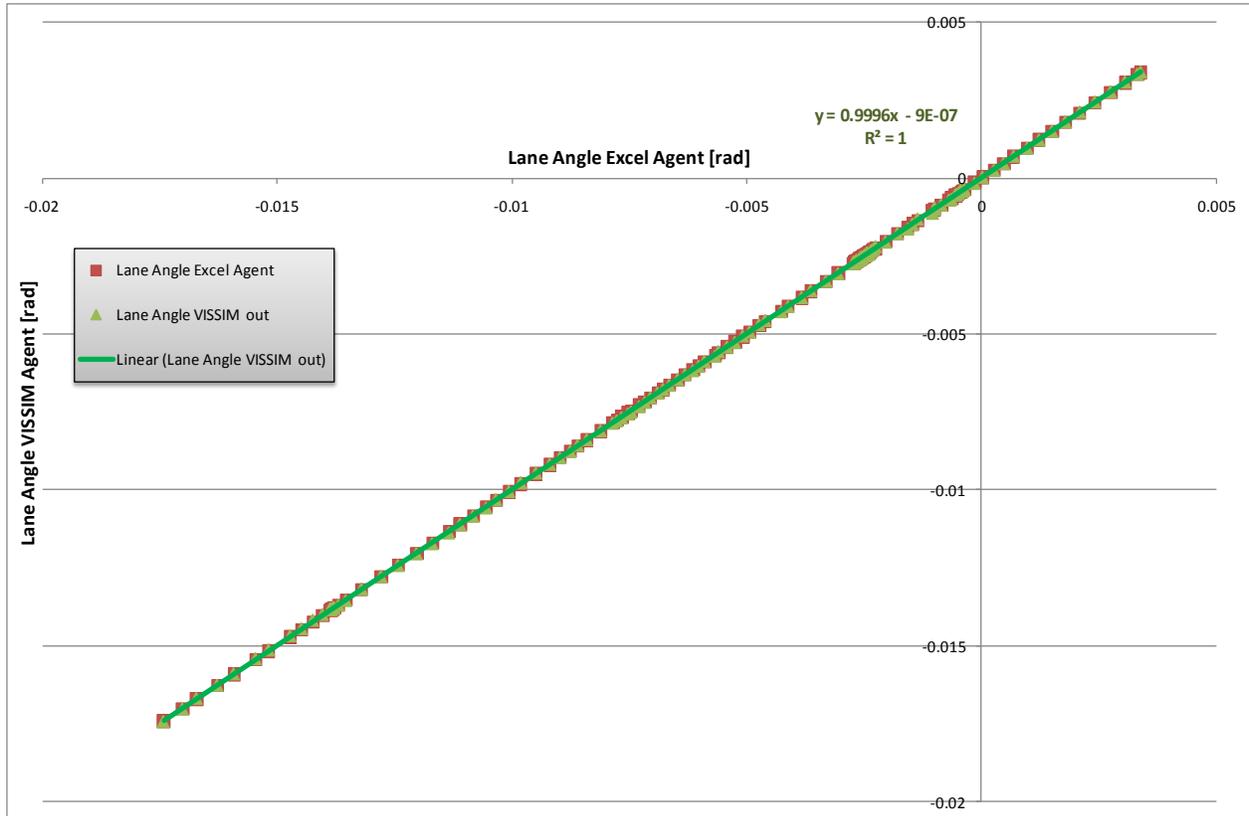


Figure 86. Graph. Lane angle scatter plot with linear regression analysis, Excel agent versus VISSIM agent.

For testing, there is another text file written to the working directory called “Test_[Name of the used parameter File].txt.” It contains all the parameters from the parameter file, several pieces of intermediate data, and the results of the model for each time step. This file can be used to compare the agent behavior in VISSIM with the agent behavior in the Excel representation of the model as well.

CHAPTER CONCLUSION

This chapter details the procedure for utilizing the new agent method in VISSIM. The series of instructions addresses how to change from the current state of practice to the new agent method. Following these instructions will result in a simulation that uses the new approach, and the results obtained from that simulation will include the behaviors that cannot be modeled by the current model but can be modeled by the new agent approach. The next chapter presents the results of the agents implemented in VISSIM versus the original Wiedemann car-following model as used by VISSIM.

CHAPTER 8. SIMULATION AND ANALYSIS

INTRODUCTION

Previous tasks have illustrated the feasibility of developing and training agents to respond to safety-critical events in a way that mimics the behavior of actual drivers observed in the naturalistic dataset. This chapter documents the results of the task with the intent to insert those agents into VISSIM and compare their performance to default drivers using the Wiedemann model for driver behavior.

Simulation Environment

The agents were tested within VISSIM version 5.30, build 8. Very simple networks were created for each of the 20 developed agents based on scenarios derived from the naturalistic data. The networks consisted of one-lane roadways with a leader and following vehicle. The leader is designed to behave in a way that creates a safety-critical event and causes the follower (the agent) to react. The duration of each simulation is 100 s. Simulation runs were conducted with and without the agents activated so that the impacts of the agent behavior could be evaluated relative to the Wiedemann model. Each of the 20 scenarios (both with and without the agent) was replicated 20 times to account for the variability in conditions.

Evaluation

The agents were evaluated from both an operational and a safety perspective. The operational analysis was conducted by collecting and analyzing the vehicle trajectory files from each run with and without the agent active. Safety was analyzed using the Surrogate Safety Assessment Model (SSAM) developed by the Federal Highway Administration. SSAM identifies conflicts in the traffic stream as modeled by VISSIM and reports a number of measures from which safety can be evaluated. The results of the operational and safety analyses are discussed in detail in the following sections. The chapter concludes with a discussion of agent scalability. A slightly more complex network was developed to evaluate the performance of the agents in situations beyond the specific conditions for which they had been trained.

OPERATIONS ASSESSMENT

The impact of the agents on vehicle dynamics and operational measures was assessed first. This section first presents a qualitative description of the observed agent behavior in VISSIM. Next it presents an analysis of objective vehicle dynamics metrics to illustrate differences in driver behavior between the agents and the Wiedemann model using the simple test network.

Agent Behavior in VISSIM

The models for the car and truck agents were tested in VISSIM and evaluated for their effects on traffic operations and mobility. An overview of the observed agent behavior on this simple one-lane VISSIM test network appears below.

Car Agents

- **Agent a:** The agent enters the network at a higher speed than the leader. Upon closing with the leader, the agent decelerates, swerves slightly, and passes (overruns) the leader.
- **Agent b:** The agent enters the network at a higher speed than the leader. Upon closing with the leader, the agent overruns the leader.
- **Agent c:** The agent enters the network at a higher speed than the leader. Upon closing with the leader, the agent decelerates significantly. Depending on the individual agent's reaction, the agent either comes to a complete stop prior to colliding with the leader or overruns the leader.
- **Agent d:** The agent enters the network at a higher speed than the leader. Upon closing with the leader, the agent decelerates significantly. Depending on the individual agent's reaction, the agent either comes to a complete stop prior to colliding with the leader or overruns the leader.
- **Agent e:** The agent enters the network at a higher speed than the leader. Upon closing with the leader, the agent swerves slightly and decelerates sharply behind the leader. The agent usually comes to a complete stop behind the leader but occasionally overruns the leader if it cannot brake in time.
- **Agent f:** The leader travels faster than the agent, and the gap between vehicles expands during simulation. No significant interaction is observed between the vehicles.
- **Agent g:** The agent enters the network at a higher speed than the leader. Several behaviors were observed when the agent and the leader interact. If the agent cannot stop in time, it sometimes overruns the leader. Otherwise, the agent decelerates quickly and swerves slightly as it brakes behind the leader. After dwelling for some time, the agent accelerates sharply until it catches up with the leader again, at which point it again decelerates and swerves off the road, sometimes by a significant margin.
- **Agent h:** The agent enters the network at a higher speed than the leader. When it catches up to the leader, it consistently overruns the leader.
- **Agent i:** Several behaviors were observed for this agent. The agent enters the network at a higher speed than the leader. When it catches up to the leader, it usually decelerates and swerves, but ultimately it overruns the leader. Sometimes the agent will catch up to the leader and maintain a safe following distance, and other times it decelerates quickly and leaves the roadway.
- **Agent j:** The agent enters the network at a higher speed than the leader. When it catches up to the leader, it consistently overruns the leader.
- **Agent k:** The agent enters the network at a higher speed than the leader. When it catches up to the leader, it consistently overruns the leader.

Truck Agents

- **Agent A:** The agent enters the network at a higher speed than the leader. When it catches up to the leader, it swerves off the right side of the road to avoid the leader, passing it in the process.
- **Agent B:** The agent enters the network at a higher speed than the leader. The agent swerves off the right side of the road and returns back to the road, coming to a complete stop perpendicular to the roadway. The distance of the roadway departure changes based on the simulation random number seed.
- **Agent C:** Several behaviors were observed for this agent. The agent enters the network at a higher speed than the leader. When it catches up to the leader, it usually decelerates and swerves, passes the agent, and returns to the roadway. However, sometimes the agent swerves and does not pass the leader, and instead continues to swerve farther and farther off the roadway, maintaining a distance of approximately 50 feet from the roadway.
- **Agent D:** The agent enters the network at a higher speed than the leader. The agent swerves slightly to the left but ultimately overruns the leader.
- **Agent E:** The agent enters the network at a higher speed than the leader. The agent swerves slightly to the left but ultimately overruns the leader.
- **Agent F:** Several behaviors were observed for this agent. The agent sometimes enters the network at a higher speed than the leader, decelerates quickly, and stops for a short period, then accelerates, catches the leader, and overruns it. Other times, the agent stops suddenly upon entering the network and after a few seconds accelerates. Occasionally it catches and overruns the leader, but other times the simulation ends before the agent catches the lead vehicle.
- **Agent G:** The agent enters the network at a higher speed than the leader. When it catches up to the leader, the agent decelerates quickly and consistently overruns the leader, eventually accelerating to its original speed.
- **Agent H:** The agent enters the network at a higher speed than the leader. The agent swerves slightly to the right and ultimately overruns the leader.
- **Agent I:** The agent enters the network at a higher speed than the leader. When it catches up to the leader, it consistently overruns the leader.

A visual inspection of the simulations demonstrates some unexpected behavior in certain agents. Some agents appear to overrun the lead vehicle in some scenarios (car agents a, b, c, e, g, h, i, j, and k and truck agents D, E, F, G, H, and I), and some agents show multiple behaviors based on the simulation seed (car agents c, d, e, g, and i and truck agents C and F).

To show the agents' behavior visually, two plots were made for each simulation run of each agent. The set of figures shows the time-space diagrams of a lead vehicle, an agent, and a following vehicle using the Wiedemann model as implemented in VISSIM. These figures show

Figure 88 and Figure 89 show plots of headway versus difference in speed between the agent and the lead vehicle. The figures show a scenario with the agent driver model and one without. In these figures, when an agent's headway is at the highest point in the chart, it indicates that the agent vehicle has passed the lead vehicle. The figures show most of the agents initially in what appears to be the "closing in" zone. As the agent approaches the lead vehicle, the agent may decelerate at the last minute, as seen in Figure 88 with Agent g, or may continue into the emergency regime with little or no reaction, as seen in Figure 89 with Agent I. Other agents appear to follow the lead vehicle, creating the circular plot line similar to Wiedemann model vehicles. This generally reflects the agent behavior described earlier, where some agents do not appear to react to the lead vehicle by slowing, but instead by swerving or driving as before until they have passed the lead vehicle, while other agents attempt to decelerate at varying rates.

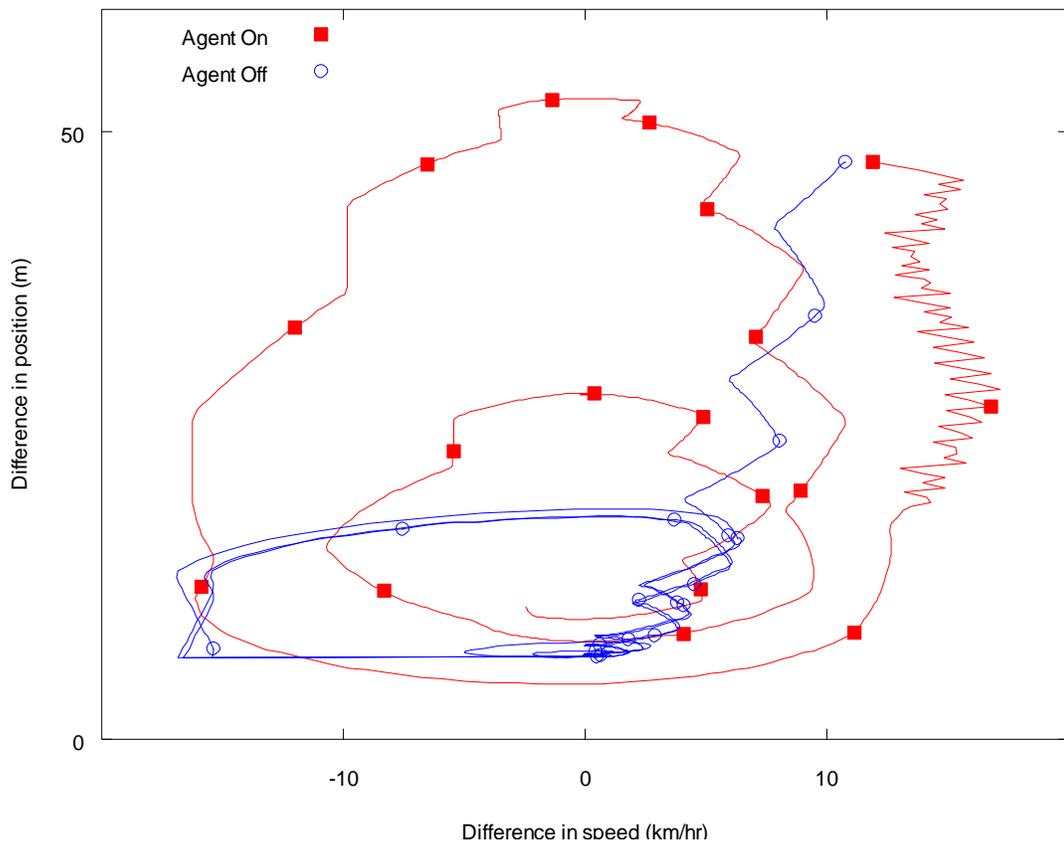


Figure 88. Chart. Wiedemann diagram of a single simulation run of Agent g.

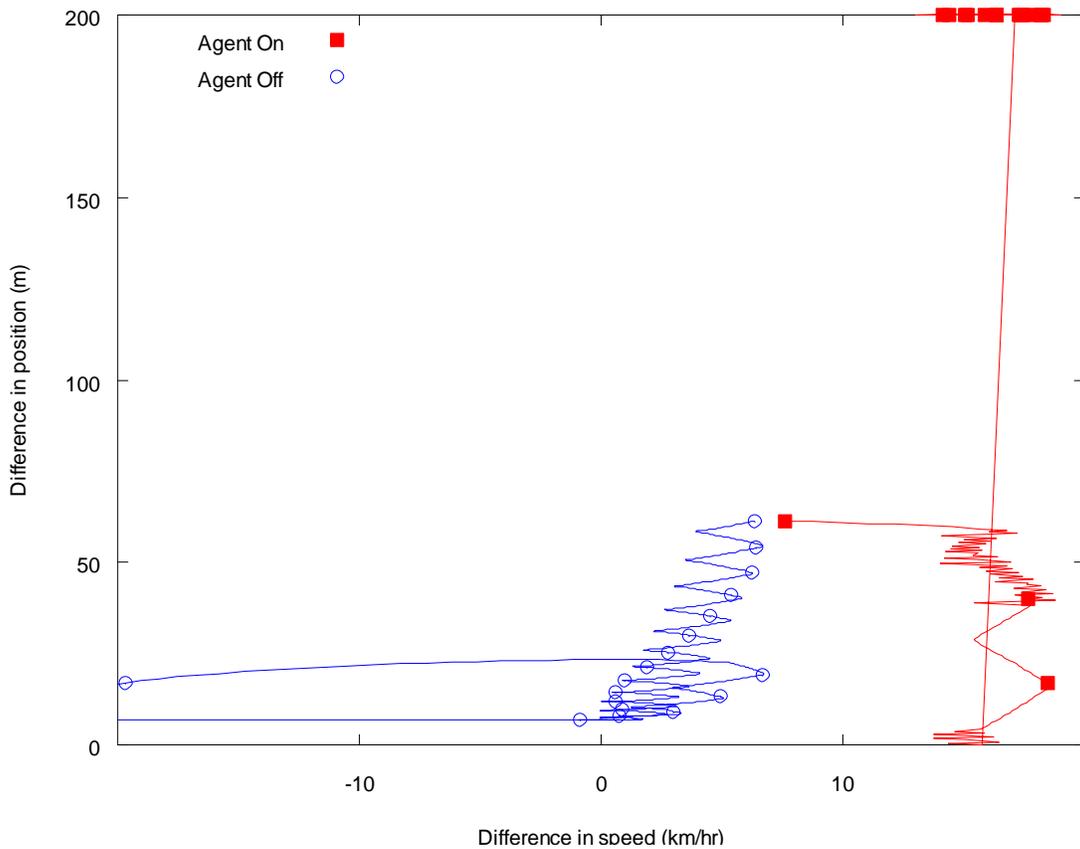


Figure 89. Chart. Wiedemann diagram of a single simulation run of Agent I.

Quantitative Assessment of Vehicle Behavior

Because the model was restricted to a single-lane, one-way network by design, and because the agents themselves could not transition from the agent-specific driving model to the general Wiedemann model, the agent could not be tested in a traditional multilane lane network with interchanges and intersections. Therefore more traditional macroscopic metrics such as travel time and delay could not be used, simply because the test networks consisted of two vehicles in a single interaction.

Several metrics were used to better understand the behavior of the agent:

- **Minimum following distance (m):** The minimum distance between the lead vehicle and the agent. A value of 0.0 indicates that the agent collides with the lead vehicle or passes it while traveling off of the roadway link.
- **Minimum acceleration (m/s^2):** The minimum acceleration rate experienced by the agent in a single run. Several agents had unusually high deceleration rates of greater than $50 m/s^2$ when first inserted. These values were considered to be arbitrary simulation error and were ignored in the analysis.

- **Average acceleration (m/s²):** The average acceleration rate of an agent over an entire simulation run.
- **Average of absolute acceleration (m/s²):** The average of the absolute value of an agent’s acceleration over an entire simulation run. High values indicate frequent and/or drastic changes in speed.
- **Average speed (km/hr):** The average speed of an agent.

Each metric was evaluated with the agent (i.e., “Agent On”) and in a base case using the Wiedemann model as implemented in VISSIM (i.e., “Agent Off”). Table 38 shows the average value for each metric over 20 simulation runs of the car agents. Table 39 shows the same information for the truck agents.

In general, most agents eventually passed the lead vehicle, as demonstrated by the minimum following distance of 0 m, while the lead vehicle was never passed in the base cases. Because a one-lane network was simulated, the agent vehicle either departed the roadway to pass the lead vehicle or drove through the lead vehicle (a collision event). Most agents (with the exceptions of agents D, E, and F) were able to eliminate the unrealistically high decelerations seen in the base cases when collision was imminent. However, because many of the agents appeared to overrun the lead vehicles, these more realistic decelerations may not in fact demonstrate more realistic behavior overall. Average acceleration was much lower in the base cases than for the agents, which is likely due to the very high decelerations used by the following vehicle to avoid collisions. Average speed was generally much lower for the base cases, which is likely a result of the base case following vehicles reducing their speeds to stay behind the lead vehicles for the bulk of the simulation. In comparison, the agents often passed the lead vehicle and continued at their desired speeds.

Table 38. Mobility performance metrics for car agents, average values over 20 simulation runs.

Agent	Agent Status	Minimum Following Distance (m)	Minimum Acceleration (m/s ²)	Average Acceleration (m/s ²)	Average Absolute Acceleration (m/s ²)	Average Speed (km/hr)
a	On	0.00	-6.46	0.00	2.62	36.68
	Off	6.27	-100.00	-0.92	1.32	24.38
	Δ	-6.27	93.54	0.92	1.30	12.31
b	On	0.00	-8.18	0.01	3.29	13.74
	Off	6.26	-100.00	-3.28	3.81	5.99
	Δ	-6.26	91.82	3.29	-0.52	7.75
c	On	2.08	-7.59	-0.44	2.51	8.56
	Off	6.26	-100.00	-3.28	3.81	5.99
	Δ	-4.18	92.41	2.84	-1.29	2.57
d	On	2.29	-9.36	0.01	2.55	5.61

Agent	Agent Status	Minimum Following Distance (m)	Minimum Acceleration (m/s ²)	Average Acceleration (m/s ²)	Average Absolute Acceleration (m/s ²)	Average Speed (km/hr)
	Off	6.26	-100.00	-3.28	3.81	5.99
	Δ	-3.98	90.64	3.29	-1.26	-0.38
e	On	4.10	-6.78	-0.01	0.81	6.67
	Off	6.26	-100.00	-3.31	3.84	5.98
	Δ	-2.17	93.22	3.30	-3.02	0.69
f	On	14.52	-8.53	0.01	3.59	3.62
	Off	6.26	-100.00	-3.02	3.57	5.93
	Δ	8.26	91.48	3.03	0.02	-2.31
g	On	4.02	-6.59	-0.03	1.20	23.06
	Off	6.26	-100.00	-1.15	1.56	19.01
	Δ	-2.25	93.41	1.12	-0.36	4.05
h	On	0.00	-7.20	0.01	2.74	32.95
	Off	6.26	-100.00	-1.15	1.56	19.01
	Δ	-6.26	92.80	1.16	1.18	13.94
i	On	2.38	-6.78	0.00	2.58	29.92
	Off	7.50	-91.82	-0.46	0.72	28.07
	Δ	-5.12	85.04	0.45	1.85	1.86
j	On	0.00	-6.46	0.00	2.62	36.68
	Off	6.27	-100.00	-0.92	1.32	24.38
	Δ	-6.27	93.54	0.92	1.30	12.31
k	On	0.00	-7.32	-0.09	2.91	19.10
	Off	6.26	-100.00	-1.27	1.64	14.15
	Δ	-6.26	92.68	1.17	1.27	4.94

Values shown in bold are significant at $\alpha = 0.05$.

Table 39. Mobility performance metrics for truck agents, average values over 20 simulation runs.

Agent	Agent Status	Minimum Following Distance (m)	Minimum Acceleration (m/s ²)	Average Acceleration (m/s ²)	Average Absolute Acceleration (m/s ²)	Average Speed (km/hr)
A	On	0.00	-6.54	0.02	2.31	71.67
	Off	6.27	-99.49	-0.49	0.85	48.28

Agent	Agent Status	Minimum Following Distance (m)	Minimum Acceleration (m/s ²)	Average Acceleration (m/s ²)	Average Absolute Acceleration (m/s ²)	Average Speed (km/hr)
	Δ	-6.27	92.95	0.51	1.46	23.38
B	On	0.00	-7.08	-2.00	2.46	13.11
	Off	13.08	-90.66	-0.33	0.69	56.58
	Δ	-13.08	83.58	-1.66	1.78	-43.48
C	On	0.82	-4.11	0.01	1.00	91.35
	Off	6.27	-98.26	-0.24	0.88	86.15
	Δ	-5.45	94.15	0.25	0.12	5.20
D	On	0.00	-21.49	0.02	3.71	18.97
	Off	6.27	-100.00	-0.97	1.24	16.40
	Δ	-6.27	78.51	0.99	2.46	2.56
E	On	0.00	-81.87	0.00	2.94	53.06
	Off	6.27	-100.00	-0.44	0.79	43.27
	Δ	-6.27	18.13	0.44	2.14	9.79
F	On	0.00	-100.00	-10.81	12.30	83.36
	Off	6.26	-100.00	-0.34	0.92	79.84
	Δ	-6.26	0.00	-10.47	11.38	3.52
G	On	0.00	-8.16	0.06	2.62	76.18
	Off	6.27	-100.00	-0.68	1.23	42.65
	Δ	-6.27	91.84	0.74	1.39	33.53
H	On	0.00	-6.73	0.05	2.44	68.52
	Off	6.26	-100.00	-0.55	1.04	57.87
	Δ	-6.26	93.27	0.60	1.40	10.65
I	On	0.00	-9.09	0.03	3.50	38.17
	Off	6.27	-96.48	-0.62	0.88	24.50
	Δ	-6.27	87.39	0.64	2.63	13.67

Values shown in bold are significant at $\alpha = 0.05$.

Summary of Operations Analysis

The qualitative and quantitative analysis on the simple linear VISSIM network revealed the following findings:

- Most agents were able to enter the emergency regime, where an agent's speed is greater than the lead vehicle's, yet headway is small. Vehicles following the Wiedemann model as

implemented in VISSIM are unable to enter the emergency regime, and instead vehicles will be forced to decelerate at unrealistic rates to avoid collisions.

- Almost all agents were able to eliminate the improbably high deceleration rates used by vehicles following the Wiedemann model as implemented in VISSIM. The agents instead swerved or ran over the leading vehicles (collision). The vast majority of the agents' deceleration rates to avoid collisions were realistic.
- Although agents experienced realistic deceleration rates, in many cases they did not behave consistently across all simulation runs. The agents' behavior indicates that although each agent is programmed to respond a certain way to a certain scenario, the definition of that scenario may be sufficiently narrow that changes in the random seed fail to re-create it. The definition of the scenario may need to be broadened or better defined so that vehicles can behave in a more consistent way. This issue can be addressed by exposing the agents to more datapoints during the training process.
- The agents were trained to emphasize safety-critical modeling, and therefore the frequency of crashes and near crashes during an agent-based episode was exaggerated. This was intentional in this project to showcase the ability of agents to simulate safety-critical events. However, for the purpose of real analysis, the agents should take into account the realistic probability of crashes and produce appropriate statistics. This should be a focus for future research.

SAFETY ASSESSMENT

A series of safety surrogate measures were analyzed to compare the behaviors of the agents and the VISSIM Wiedemann car-following model. This provided an indication as to whether the agents were creating significantly more safety-critical events than the base VISSIM car-following model. The Surrogate Safety Assessment Model (SSAM), version 2.1.6, was used to examine the influence of the agents and the Wiedemann model on safety surrogate measures. The SSAM can post-process vehicle trajectory files generated by VISSIM and provides a mechanism to compare estimated safety surrogate measures with and without the various agents.

The SSAM identifies conflict events based on time-to-collision (TTC) thresholds. TTC is defined as the elapsed time before two vehicles collide if they maintain their current speed and heading. In this analysis, conflicts were identified using two TTC thresholds: 1.5 s and 2.5 s. A TTC of 2.5 s is a relatively high threshold for triggering conflicts, but it was set high to ensure that some conflict events would be triggered when the Wiedemann car-following model was active. The 2.5-s value was based on the assumed perception-reaction time for stopping sight distance in the American Association of State Highway and Transportation Officials' *Policy on the Geometric Design of Highways and Streets*.

When the conflict TTC threshold is exceeded, the SSAM generates a series of safety surrogate measures for each conflict event. Specific measures examined in this research include the following:

- **TTC** is the minimum time to collision between the agent and leading vehicle during the conflict. If a conflict occurs and has a TTC of 0, it means that a collision will occur.

- **PET** is the minimum post encroachment time observed during the conflict. This represents the elapsed time between when the leading vehicle occupies a point and when the agent arrives at the same position. A PET of zero indicates a collision has occurred.
- **MaxS** is the maximum speed of either vehicle throughout the conflict.
- **DeltaS** is the difference in speeds observed at the time step when the minimum TTC occurs. If the minimum TTC equals zero, this would be the relative difference in speeds at collision.
- **DR** is the initial deceleration rate of the following vehicle (agent) during the conflict. If the following vehicle brakes, this is the first negative acceleration value observed during the conflict. If the follower does not brake, then this is the lowest acceleration value observed during the conflict.
- **MaxD** is the maximum deceleration of the following vehicle (agent). Negative values indicate deceleration, and positive values indicate acceleration.
- **MaxDeltaV** is the maximum change in velocity between when a conflict first begins and immediately post-crash, assuming a hypothetical inelastic crash between vehicles. This is a surrogate for conflict severity.

All of these measures are tabulated when a conflict has been identified. In other words, if a conflict does not occur, then zeros are not included in the averages.

Each agent scenario was examined with both the agent active and with the agent inactive (Wiedemann car-following active). Twenty replications were conducted for each scenario and each agent status. This resulted in a total of 800 simulation trajectories (20 agents \times 20 replications \times 2 levels of agent status [agent active or Wiedemann active]) that were evaluated using SSAM. These were evaluated using the 1.5-s and 2.5-s TTC values for conflict identification.

General Observations on Agent Behavior and Impact on Safety Surrogates

The reactions of agents were observed during the course of each simulation to identify responses that could impact the safety surrogate measures being evaluated. First, even though the simulation duration was only 100 s, several agents would produce multiple safety-critical events over the course of the simulation. For example, an agent sometimes approached the leader and braked sharply, coming to a complete stop. After dwelling at a stop for a few seconds, the agent accelerated sharply until it caught up to the leader again. At that point, it braked sharply a second time. This sort of behavior could produce multiple conflict events in one simulation replication, but the behavior of the agent was not necessarily realistic because it just repeated the same unsafe behavior multiple times.

A second behavior that was commonly seen in the agent simulations was the agent overrunning (traveling completely over) the leading vehicle. In this case, the safety surrogates would show a value of 0 s for the TTC and PET because a collision would occur. After passing over the leader, however, both the leader and agent proceeded as if no collision had occurred. The behavior after the collision was not necessarily realistic in many cases.

Finally, some agents would swerve around the lead vehicle, sometimes actually leaving the roadway link. While this behavior is certainly a conflict, it may not show up strongly in some of the safety surrogate measures. When a vehicle swerves around the leader, it may not have a short enough TTC to trigger a conflict identification because it has moved onto a path where it will not conflict with the lead vehicle. This would mean that critical safety surrogate measures for these agents may be underreported.

The safety surrogates provided by SSAM provide some general indications of performance of the agents, but there are definitely cases where they may not fully represent the agent behavior. As a result, many of the safety surrogates should be viewed for relative changes between the agent and the Wiedemann model, as opposed to absolute values of each measure.

Evaluation of Car Agents

Table 40 summarizes the safety surrogate measures for the car agents, using a 1.5-s TTC to define a conflict. In addition to the surrogate measures discussed earlier, the average number of conflicts per run is also listed. All conflicts evaluated were rear-end conflicts because the agents were developed based on rear-end conflict cases. No other conflict types were produced for any agent. The Wiedemann car-following model (noted as Agent Off in Table 40), did not produce any conflicts when the TTC threshold was set to 1.5 s. This shows that the Wiedemann model does not permit close following behaviors by default, and thus uses the assumption of safe driving behavior.

With the exception of Agent f, all agents produced conflicts at least occasionally. Agent f entered the network at a lower speed than the leading vehicle, meaning the agent never caught up to the leader to cause a conflict. The remaining agents produced a mean number of conflicts of between 0.55 and 1.75 (11 and 35 conflicts per 20 runs, respectively). Mean TTCs and PETs were 0 s for 5 of 11 car agents: a, d, h, i, and j. This means that all of these agents produced collisions between the agent and the leading vehicle.

Table 40. Safety surrogate measure comparison using 1.5-s TTC – car agents.

Agent	Agent Status	Mean # of Conflicts	Mean TTC (s)	Mean PET (s)	Mean MaxS (m/s)	Mean DeltaS (m/s)	Mean DR (m/s ²)	Mean Max D (m/s ²)	Mean MaxDeltaV (m/s)
a	On	1.0	0	0	10.44	3.95	-0.13	-0.13	1.98
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
b	On	1.75	0.62	0.43	4.11	2.42	-2.58	-3.01	1.21
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.75	n/a	n/a	n/a	n/a	n/a	n/a	n/a
c	On	0.80	0.23	0.12	3.93	2.31	-0.24	-1.24	1.16
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	0.80	n/a	n/a	n/a	n/a	n/a	n/a	n/a
d	On	0.65	0	0	2.36	0.72	-0.05	-0.11	0.36
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	0.65	n/a	n/a	n/a	n/a	n/a	n/a	n/a
e	On	0.55	0.46	0.11	2.55	0.77	0.16	-0.45	0.39
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	0.55	n/a	n/a	n/a	n/a	n/a	n/a	n/a
f	On	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
g	On	0.60	0.46	0.09	8.11	3.21	-0.72	-1.72	1.61
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	0.60	n/a	n/a	n/a	n/a	n/a	n/a	n/a
h	On	1.0	0	0	9.44	4.13	-0.15	-0.16	2.03
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
i	On	0.85	0	0	8.67	1.15	-0.38	-2.83	0.57
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	0.85	n/a	n/a	n/a	n/a	n/a	n/a	n/a
j	On	1.0	0	0	10.44	3.95	-0.13	-0.13	1.98
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
k	On	1.60	0.60	0.24	6.04	2.25	-2.50	-3.55	1.12
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.60	n/a	n/a	n/a	n/a	n/a	n/a	n/a

n/a: Measure or comparison not generated because no conflicts were detected.

Table 41 summarizes the safety surrogate measures for the car agents using a TTC threshold of 2.5 s to define a conflict. When the TTC threshold was raised to 2.5 s, the Wiedemann model sometimes produced conflicts for some agent scenarios. The scenarios for Agents b, c, d, e, f, and k (6 of 11 scenarios) then had some conflicts detected when the Wiedemann car-following model was used. Once again, all agents except for Agent f produced some conflicts, with the mean number of conflicts ranging from 0.7 to 1.95 (14 to 39 conflicts per 20 runs). Mean TTC and PET values of the agents increased because conflicts at higher TTCs were then included in the summary.

For the five cases where both the agent and the Wiedemann model produced conflicts, it was possible to examine differences between the surrogate measures using t-tests. Values shown in bold in Table 41 were significant at $\alpha = 0.05$. As expected, the agents produced more safety-

critical conflicts, and the surrogate measures generally showed that the conflicts were more severe than when the Wiedemann model was used.

Table 41. Safety surrogate measure comparison using 2.5-s TTC – car agents.

Agent	Agent Status	Mean # of Conflicts	Mean TTC (s)	Mean PET (s)	Mean MaxS (m/s)	Mean DeltaS (m/s)	Mean DR (m/s ²)	Mean Max D (m/s ²)	Mean MaxDeltaV (m/s)
a	On	1.0	0	0	10.45	4.06	-0.08	-0.15	2.03
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
b	On	1.95	1.21	0.82	4.10	2.41	-2.39	-3.57	1.20
	Off	0.20	2.20	1.47	2.73	1.27	-0.12	-4.49	0.63
	Δ	1.75	-0.995	-0.657	1.373	1.141	-2.277	0.923	0.570
c	On	1.45	1.37	0.98	4.10	2.46	-3.20	-4.29	1.23
	Off	0.20	2.20	1.47	2.73	1.27	-0.12	-4.49	0.63
	Δ	1.25	-0.828	-0.499	1.370	1.188	-3.088	0.205	0.594
d	On	0.70	0.18	0.04	2.39	0.72	-0.03	-0.11	0.36
	Off	0.20	2.20	1.47	2.73	1.27	-0.12	-4.49	0.63
	Δ	0.50	-2.021	-1.432	-0.335	-0.552	0.091	4.380	-0.276
e	On	0.80	1.12	0.31	2.53	1.02	0.13	-0.67	0.51
	Off	0.20	2.20	1.47	2.73	1.27	-0.12	-4.49	0.63
	Δ	0.60	-1.081	-1.169	-0.196	-0.251	0.249	3.819	-0.125
f	On	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Off	0.30	2.25	1.62	2.43	1.08	-0.21	-18.70	0.54
	Δ	-0.30	n/a	n/a	n/a	n/a	n/a	n/a	n/a
g	On	0.80	0.88	0.21	8.41	3.68	-0.48	-1.70	1.84
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	0.80	n/a	n/a	n/a	n/a	n/a	n/a	n/a
h	On	1.05	0.22	0.13	9.44	4.38	-1.35	-1.80	2.19
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.05	n/a	n/a	n/a	n/a	n/a	n/a	n/a
i	On	0.95	0.22	0.03	8.70	1.22	-0.38	-2.51	0.61
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	0.95	n/a	n/a	n/a	n/a	n/a	n/a	n/a
j	On	1.0	0	0	10.45	4.06	-0.08	-0.15	2.03
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
k	On	1.80	1.11	0.45	6.06	2.30	-4.16	-5.86	1.15
	Off	0.25	2.20	0.66	4.11	1.22	-0.33	-0.80	0.61
	Δ	1.55	-1.092	-0.207	1.943	1.087	-3.828	-5.066	0.543

n/a: MOE or comparison not generated since no conflicts were detected

Values shown in bold are significant at $\alpha = 0.05$.

Evaluation of Truck Agents

Table 42 summarizes the evaluation of the truck agent scenarios when the TTC threshold for conflicts was set to 1.5 s. Like the car agent scenarios, the Wiedemann car-following model never produced any conflicts when the TTC threshold was set to 1.5 s. When the truck agents were active, seven of nine agents produced conflicts at least occasionally. For those agents that produced conflicts, the mean number of conflicts ranged from 0.8 to 1.60 conflicts per run (16 to

32 conflicts per 20 runs). Agents D, E, F, G, H, and I all had a mean TTC and PET of 0, meaning that they produced collisions in each conflict generated.

In the case of Agents A and B, no conflicts were produced. Agent A consistently swerved around the lead vehicle by a significant margin, leaving the network and re-entering the link in front of the leader. While this was definitely a significant safety-critical event, the agent never triggered the 1.5-s TTC threshold to define a conflict, so no safety surrogate measures were generated by SSAM. A similar situation happened with Agent B. In this case, the vehicle decelerated and swerved off of the roadway link before it could get into a conflict with the lead vehicle. Thus, even though a conflict event was not triggered for either agent, the agent behavior was definitely safety critical.

Table 42. Safety surrogate measure comparison using 1.5-s TTC – truck agents.

Agent	Agent Status	Mean # of Conflicts	Mean TTC (s)	Mean PET (s)	Mean MaxS (m/s)	Mean DeltaS (m/s)	Mean DR (m/s ²)	Mean Max D (m/s ²)	Mean MaxDeltaV (m/s)
A	On	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
B	On	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
C	On	0.80	0.55	0.06	30.07	6.46	-0.62	-3.18	3.23
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	0.80	n/a	n/a	n/a	n/a	n/a	n/a	n/a
D	On	1.10	0	0	5.95	1.82	0.48	0.48	0.91
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.10	n/a	n/a	n/a	n/a	n/a	n/a	n/a
E	On	1.60	0	0	16.31	3.79	1.45	1.45	1.90
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.60	n/a	n/a	n/a	n/a	n/a	n/a	n/a
F	On	1.05	0	0	27.91	8.13	-1.16	-2.55	4.07
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.05	n/a	n/a	n/a	n/a	n/a	n/a	n/a
G	On	1.0	0	0	21.52	9.41	-2.77	-5.09	4.71
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
H	On	1.05	0.03	0	19.51	1.93	-0.03	-0.28	0.96
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.05	n/a	n/a	n/a	n/a	n/a	n/a	n/a
I	On	1.0	0	0	10.96	4.52	-0.01	-0.01	2.26
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.0	n/a	n/a	n/a	n/a	n/a	n/a	n/a

n/a: MOE or comparison not generated because no conflicts were detected.

Table 43 summarizes the safety surrogate measures for the truck agents when the TTC threshold for a conflict was set to 2.5 s. Once again, Agents A and B did not produce any conflicts because they left the roadway and did not interact with the leader. The remaining agents produced an average number of rear-end conflicts ranging from 0.95 to 1.60 (19 to 32 conflicts per 20 runs). Using the 2.5-s TTC threshold to define a conflict also created conflicts when the Wiedemann car-following model was used in scenarios for Agents A, B, D, F, and H (5 of 9 cases).

Agents D, F, and H were the only cases where both the agent and the Wiedemann model produced conflicts. The t-tests in those cases showed that the agents tended to produce significantly more conflicts than when the Wiedemann model was used, and the conflicts produced were more severe. This was as expected.

Table 43. Safety surrogate measure comparison using 2.5-s TTC – truck agents.

Agent	Agent Status	Mean # of Conflicts	Mean TTC (s)	Mean PET (s)	Mean MaxS (m/s)	Mean DeltaS (m/s)	Mean DR (m/s ²)	Mean Max D (m/s ²)	Mean MaxDeltaV (m/s)
A	On	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Off	0.10	2.25	0.20	12.18	1.16	-0.49	-1.05	0.58
	Δ	-0.10	n/a	n/a	n/a	n/a	n/a	n/a	n/a
B	On	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Off	0.10	2.50	0.20	14.34	0.90	-0.58	-0.78	0.45
	Δ	-0.10	n/a	n/a	n/a	n/a	n/a	n/a	n/a
C	On	0.95	0.78	0.11	30.19	6.27	-1.63	-3.57	3.14
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	0.95	n/a	n/a	n/a	n/a	n/a	n/a	n/a
D	On	1.30	0.38	0.18	5.89	1.78	-0.39	-0.54	0.89
	Off	0.10	2.40	0.60	4.29	0.99	-0.12	-2.76	0.49
	Δ	1.20	-2.023	-0.423	1.594	0.793	-0.277	2.220	0.396
E	On	1.60	0	0	16.31	3.51	1.44	1.44	1.75
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.60	n/a	n/a	n/a	n/a	n/a	n/a	n/a
F	On	1.05	0	0	28.18	8.35	-1.08	-2.64	4.18
	Off	0.15	2.47	0.20	20.45	1.55	-0.53	-0.56	0.77
	Δ	0.90	-2.467	-0.20	7.730	6.808	-0.554	-2.077	3.404
G	On	1.0	0	0	21.52	8.62	-2.45	-5.20	4.31
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
H	On	1.05	0.03	0	19.56	1.97	0.03	-0.30	0.98
	Off	0.10	2.35	0.30	16.34	2.09	-0.64	-0.72	1.05
	Δ	0.95	-2.317	-0.295	3.213	-0.122	0.673	0.418	-0.061
I	On	1.0	0	0	10.96	4.53	0	-0.01	2.26
	Off	0	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Δ	1.0	n/a	n/a	n/a	n/a	n/a	n/a	n/a

n/a: MOE or comparison not generated because no conflicts were detected.

Values shown in bold are significant at $\alpha = 0.05$.

Summary of Safety Surrogate Analysis

The safety surrogate measure analysis on the simple linear VISSIM network revealed the following findings:

- As expected, the agents produced many more safety-critical events than when VISSIM used the Wiedemann car-following model to control the following vehicle. The Wiedemann car-following model never produced any conflicts when the TTC threshold for a conflict was set at 1.5 s, regardless of the scenario being evaluated.

- Several of the agents were not well suited to analysis using SSAM. Agents that departed the roadway by a significant margin were not properly classified as having caused safety-critical events.
- When the TTC threshold for conflicts was increased to 2.5 s, the Wiedemann model sometimes produced conflicts. However, the number and severity of conflicts with the Wiedemann model were not as severe as when the agent was active.
- Reactions of the agents following the safety-critical event were sometimes not realistic. Agents would repeat behaviors that created safety-critical events multiple times, or would proceed following a collision as if nothing had happened. Before the agents can be implemented in a more complex model, behaviors following critical events will need to be re-examined for the agents. In effect, noncollision events would need to reset to the traditional Wiedemann model, and collision events would need to be modeled by having both vehicles come to a stop after the event.

EVALUATION OF AGENTS ON MORE COMPLEX NETWORKS

A simple VISSIM test network was developed to evaluate agent performance on a network more complex than the single-lane, unidirectional link evaluated earlier. The intent of this analysis was to determine the potential impact of agent behaviors on macroscopic traffic characteristics, for example by examining the impact of safety-critical events on traffic flow and travel time reliability. The test network consisted of a 4-mi stretch of freeway with two interchanges. One of the interchanges in the example network is shown in Figure 90 below. Traffic volumes were set so that transient congestion formed at freeway on-ramps and there was turbulence due to lane changing at the off-ramps. The intent was to evaluate the impact of different agent types and penetrations on operational measures on this network. Since the ramp termini were near capacity using the Wiedemann car-following model, the expectation was that the agents would create more significant congestion on the network.

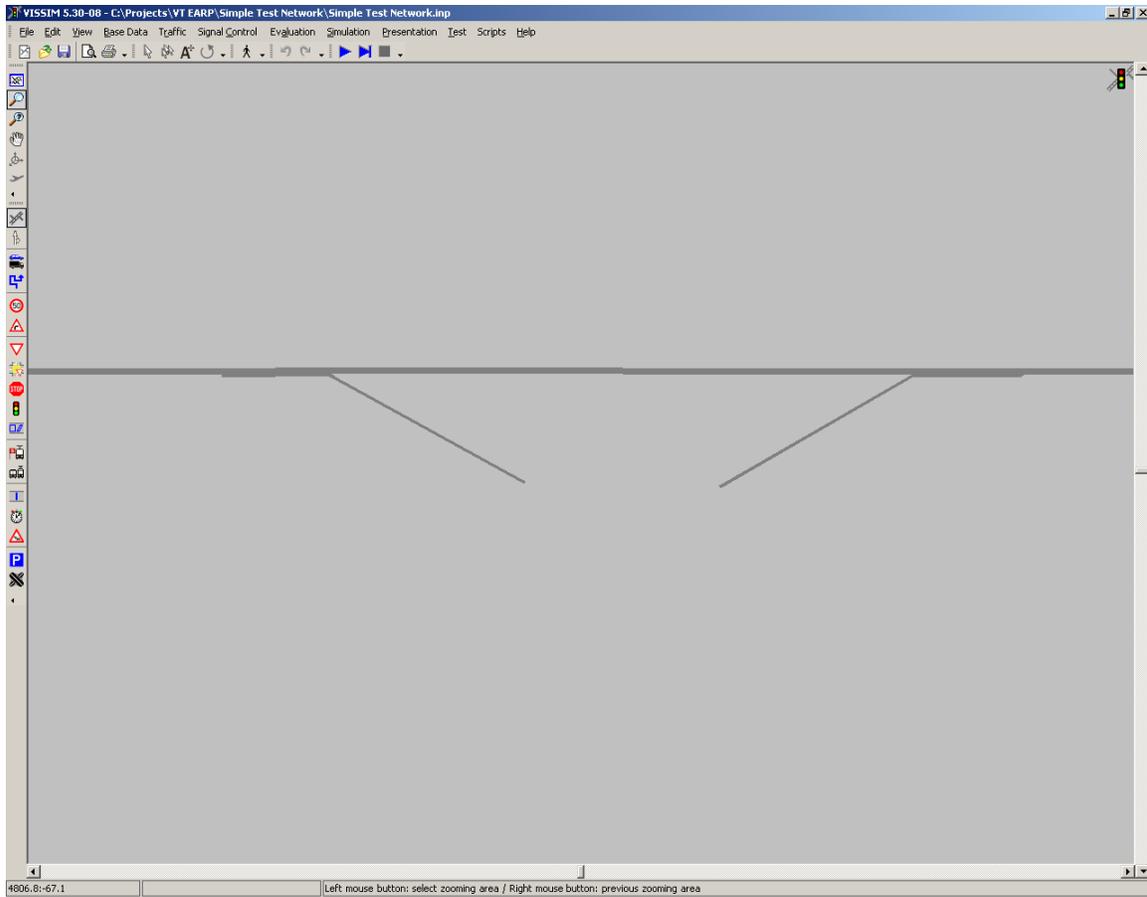


Figure 90. Screenshot. Example of interchange on test network.

Initial evaluations on this network revealed several problems with directly implementing the agents. First, many of the problems that were manifested on the simple linear network continued to occur on this network. Agents would crash with leading vehicles with much higher frequency than realistically expected, and the lead vehicles were not affected by the actions of the subject vehicles. The net result of these issues was that traffic operations on this network did not appear to be realistic because lead vehicle interactions with the agents were not reasonable.

The effect was that many more safety-critical events were occurring in the simulation than could be reasonably expected. As a result, a detailed evaluation of the agents on this network could not be conducted.

CHAPTER CONCLUSION

This chapter shows the results obtained from implementing the new agent approach into VISSIM as compared to the state of practice. The results showed that the agents would adopt more realistic decelerations than the Wiedemann model.

Before the agents are ready to be implemented in a production environment in VISSIM, several issues need to be resolved:

- The areas of application (speeds, geometric conditions) where specific agents should be used needs to be defined. Before being used in a production environment, the selection of appropriate agents for a specific condition should be automated.
- As a future research item, there is a need to estimate the frequency with which certain agent behaviors occur in the driving population to ensure that they are modeled at appropriate levels.

Once these changes are made, the inclusion of agents in VISSIM would offer a powerful opportunity to examine safety implications of different designs, as well as to better assess the impacts of incidents and other unpredictable driver behavior on travel time reliability. The initial testing of the agents in VISSIM shows that the use of agents is viable conceptually, but additional refinement is needed before they can be used outside of a research environment focused on specific events.

CHAPTER 9. CONCLUSIONS AND RECOMMENDATIONS

CONCLUSIONS

This report has documented the research performed to model driver behavior in traffic under naturalistic driving data. VTTI naturalistic driving data was used in this research due to its unique capabilities to capture the behavior of drivers when they are not in a test environment. Both safety-critical and car-following periods were used for training agents to mimic the range of behaviors of drivers.

The naturalistic data analysis allowed for the discovery of new thresholds for the Wiedemann model: a passing threshold and a hook following threshold. These two new thresholds increase the usability and accuracy of the Wiedemann model when natural or naturalistic driver behaviors are concerned because these thresholds were developed based on observations made in the naturalistic data. A hybrid model was developed, combining the strengths of both the GHR model and the Wiedemann model to yield a more accurate and highly flexible model. The comparison of different car-following models resulted in finding that different models have the least error for different drivers, but the velocity difference model and the Wiedemann model both perform adequately across all the drivers tested.

To model safety-critical events in addition to normal events, NFACRL, an agent-based artificial intelligence machine-learning technique, was used to model driving behavior. The naturalistic driving database was used to train and validate driver agents. The advantages of NFACRL lie mainly in its ability to simulate heterogeneous vehicle actions in complicated traffic environments. It is worth mentioning that this research was an attempt to apply reinforcement learning techniques to solving high-dimensional state problems and continuous-action simulation problems simultaneously in transportation research, especially in traffic flow theory. From the perspective of microscopic traffic behavior modeling, the proposed methodology was able to simulate lateral action, which brings new insights to the modeling of driver maneuvering behavior during safety-critical events. The proposed methodology also simulated events from different drivers and proved behavior heterogeneities. Robust agent activation techniques were also developed using discriminant analysis.

The developed agents were implemented in the VISSIM simulation platform and were evaluated by comparing the behavior of vehicles with and without agent activation. The results showed very close resemblance of the behavior of agents and driver data. Prototype agents were developed and submitted with this research report.

RESEARCH CONTRIBUTION

The research conducted in this project can be considered as a foundation for agent-based modeling and simulation based on naturalistic data. One of the major contributions of this research is the developed integrated framework for safety and operation analysis. The research has contributed the following:

1. Improvement of car-following models during normal driving episodes, namely the development of the hybrid Wiedemann-GHR model and the addition of the new pass and hook thresholds.
2. Development of the NFACRL agent-based model for safety-critical events, and training of 20 agents to mimic 20 different drivers extracted from the naturalistic database.
3. Development of agent-based simulation components integrated with the VISSIM simulation package through its driver model dll.
4. Development of a robust activation mechanism for agent-based simulation based on discriminant analysis and implementing the activation mechanism in a GUI.

RECOMMENDATION FOR FUTURE RESEARCH

There are several suggestions for future research that can be categorized under theoretical, applied, and related efforts. On the theoretical aspect, the next step of this research is the extension of the NFACRL framework to simulate other traffic behavior, such as lane-changing behavior and merging behavior in the upstream and downstream of ramps. It would be interesting to model individual driver behavior and the decision-making process under these traffic conditions. Training parameters, such as speed, memory, and scaling factor, were fixed for all the events in this research, but it would be interesting to see how NFACRL performs under different combination of these factor sets. As with most artificial intelligence methods, NFACRL is a heuristic approach, so there is no guarantee that training results would be optimal. Therefore, training parameters play an important role in guiding NFACRL to find a near-optimal result. We also found that agent performance is very sensitive to the driver-dependent training parameters, discrete action sets, and state bounds. Theoretically, it would be better if these parameters were preoptimized before training. However, because the only way to test agent performance is through NFACRL training, parameter optimization and training form a cycle, and computation time would exponentially increase. Our approach set the parameters based on statistical quartiles, and this yielded good results. Future research should focus on the optimization methodology for training parameters and will improve agent performance.

On the applied aspect, our initial evaluations on events beyond what the agents were trained for revealed that agents can behave erratically in those cases. It is therefore recommended to train agents using more detailed naturalistic datasets (e.g., SHRP II naturalistic datasets). In addition, it would also be interesting to examine the behavior of agents if they were first trained to mimic the Wiedemann model behavior, for example, then overlaid with naturalistic data training. More importantly, research should be conducted to evaluate the potential of using agent-based simulation to examine different design alternatives based on safety and operation factors. The evaluation effort should examine unique or problematic geometry, network design characteristics, and traffic conditions that necessitate both operational and safety analysis (e.g., different weather conditions) and compare the results to existing and/or known outcomes.

Finally, on the related research aspect, the research conducted in this report has shown the importance of two key future research issues:

1. Adaptability of agents in real time: agents could be programmed to adapt during the simulation itself.
2. Human factor issues related to warning individual drivers about a change in their driving behavior that might lead to a safety-critical event.

REFERENCES

1. Treiber, M., A. Kesting, et al. (2006). "Delays, Inaccuracies and Anticipation in Microscopic Traffic Models." 28.
2. Gipps, P. G. (1981). "A Behavioural Car-Following Model for Computer Simulation." Transportation Research Board **15B**(2-C): 7.
3. Olstam, J. J. and A. Tapani (2004). Comparison of Car-following models, VTI Meddelande 990A, Swedish National Road Transport Research Institute. **Project Code 40503 and 40485, 2004.**
4. Ossen, S. and S. P. Hoogendoorn (2008). "Validity of trajectory-based calibration approach of car-following models in presence of measurement errors." Transportation Research Record(2088): 117-125.
5. Kesting, A. and M. Treiber (2008). "Calibrating Car-Following Models using Trajectory Data: Methodological Study." 17.
6. Herrey, E. M. J. and H. Herrey (1945). "Principles of Physics Applied to Traffic Movements and Road Conditions." American Journal of Physics **13**(1): 1-14.
7. Newell, G. F. (2002). "A simplified car-following theory: a lower order model." Transportation Research Part B: Methodological **36**(3): 195-205.
8. Pipes, L. A. (1953). "Operational analysis of traffic dynamics." Journal of Applied Physics **24**(3): 274-281.
9. Chandler, R. E., R. Herman, et al. (1958). "Traffic Dynamics: Studies in Car Following." Operations Research **6**(2): 165-184.
10. Herman, R. and R. B. Potts (1959). "Single lane traffic flow theory and experiment " Proceedings of the Symposium on the Theory of Traffic Flow.
11. Gazis, D. C., R. Herman, et al. (1959). "Car-Following Theory of Steady-State Traffic Flow." Operations Research **7**(4): 499-505.
12. Gazis, D. C., R. Herman, et al. (1961). "Nonlinear Follow-The-Leader Models of Traffic Flow." Operations Research **9**(4): 545-567.
13. Forbes, T. W. (1963). "Human Factor consideration in traffic flow theory." Highway Research Board Record **15. HRB**: 60.
14. Gazis, D. C., R. Herman, et al. (1959). "car-following theory of steady state flow." Operations Research **7**(4).
15. Wiedemann, R. (1974). "Simulation des Strassenverkehrsflusses."
16. Leutzbach, W. (1988). "Introduction to the theory of traffic flow."
17. Wiedemann, R. and U. Reiter (1992). "Microscopic Traffic Simulation, the simulation."
18. Fritzsche, H.-T. (1994). "A model for traffic simulation simulation " Transportation Engineering Contrinution **5.**
19. Nagel, K. and M. Schreckenberg (1992). "A Cellular Automaton Model for Freeway Traffic." Journal of Physics I **2.**
20. Lan, L. W., C. Yu-Chiun, et al. (2009). "A refined cellular automaton model to rectify impractical vehicular movement behavior." Physica A **388**(Copyright 2010, The Institution of Engineering and Technology): 3917-3930.
21. Hamdar, S. H. and H. S. Mahmassani (2008). "From existing accident-free car-following models to colliding vehicles exploration and assessment." Transportation Research Record(2088): 45-56.

22. Krauss, S., P. Wagner, et al. (1997). "Metastable states in a microscopic model of traffic flow." Physical Review E. Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics **55**(Compendex): 5597-5597.
23. Treiber, M., A. Hennecke, et al. (2000). "Congested traffic states in empirical observations and microscopic simulations." Physical Review E - Statistical Physics, Plasmas, Fluids, and Related Interdisciplinary Topics **62**(Compendex): 1805-1824.
24. Jiang, R., Q. Wu, et al. (2001). "Full velocity difference model for a car-following theory." Physical Review E - Statistical, Nonlinear, and Soft Matter Physics **64**(Compendex): 017101/017101-017101/017104.
25. Panwai, S. and H. Dia (2005). "Comparative evaluation of microscopic car-following behavior." IEEE TRANSACTIONS ON INTELLIGENT TRANSPORTATION SYSTEMS **6**(3): 314-325.
26. Siuhi, S. and M. Kaseko (2010). "PARAMETRIC STUDY OF STIMULUS-RESPONSE BEHAVIOR FOR CAR-FOLLOWING MODELS." TRB 2010 Annual Meeting CD-ROM.
27. Osaki, H. (1993). "Reaction and anticipation in the Car-Following Behavior." In Proceedings of the 12 International Symposium on the Theory of Traffic Flow and Transportation.
28. Subranmanian, H. (1996). Estimation of Car-Following Models. *Masters Thesis*. Massachusetts Institute of Technology.
29. Ahmed, K. I. (1999). Modeling Drivers' Acceleration and Lane Changing Behavior. *PhD Dissertation*. Massachusetts Institute of Technology.
30. Ossen, S. and S. P. Hoogendoorn (2004). Car-following Behavior Analysis from Microscopic Trajectory Data
31. Schultz, G. G. and L. R. Rilett (2004). "Analysis of Distribution and Calibration of Car-Following Sensitivity Parameters in Microscopic Traffic Simulation Models." Transportation Research Record **1876**: 11.
32. Brackstone, M. (2004). Driver Psychological Types and Car Following: Is There a Correlation? Results of a Pilot Study.
33. Menneni, S., C. Sun, et al. (2008). "An Integrated Microscopic and Macroscopic Calibration for Psycho-Physical Car Following Models " TRB 2009 Annual Meeting CD-ROM: 17.
34. Brockfeld, E., R. D. Kuhne, et al. (2004). "Calibration and Validation of Microscopic Traffic Flow Models." Transportation Research Record **1876**: 9.
35. Hoogendoorn, S. and R. Hoogendoorn (2010). "A Generic Calibration Framework for Joint Estimation of Car Following Models using Microscopic Data." TRB 2010 Annual Meeting CD-ROM: 17.
36. Punzo, V. and A. Tripodi (2007). "Steady-state solutions and multiclass calibration of gipps microscopic traffic flow model." Transportation Research Record(1999): 104-114.
37. Rakha, H., C. C. Pecker, et al. (2007). "Calibration Procedure for Gipps Car-Following Model." Transportation Research Record **1999**: 13.
38. Hamdar, S. H., M. Treiber, et al. (2008). "Modeling driver behavior as sequential risk-taking task." Transportation Research Record(Compendex): 208-217.
39. Hamdar, S. H., M. Treiber, et al. (2009). "Calibration of a Stochastic Car-Following Model Using Trajectory Data: Exploration and Model Properties." TRB 2009 Annual Meeting CD-ROM: 19.
40. Chiabaut, N., L. Leclercq, et al. (2010). "Analysis of driving behavior heterogeneity in experimental NG-Sim data set." TRB 2010 Annual Meeting CD-ROM.

41. Ranjitkar, P. and T. Nakatsuji (2010). "A TRAJECTORY BASED ANALYSIS OF DRIVERS' RESPONSE IN CAR FOLLOWING SITUATIONS." TRB 2010 Annual Meeting CD-ROM: 21.
42. Hamdar, S. H. and H. S. Mahmassani (2010). "Individual Variations versus Collective Traffic Patterns: Heterogeneity Effect in a Risk-Taking Environment." TRB 2010 Annual Meeting CD-ROM: 20.
43. Xin, W., J. Hourdos, et al. (2008). "The less-than-perfect driver: A model of collision-inclusive car-following behavior." Transportation Research Record(Compendex): 126-137.
44. Chatterjee, I. and G. A. Davis (2010). "Bayesian Trajectory-Based Reconstruction of Rear-Ending Events Using Naturalistic Driving Data." TRB 2010 Annual Meeting CD-ROM: 18.
45. Dingus, T. A., S. G. Klauer, et al. (2006). The 100-car naturalistic Driving Study; Phase II - Results of the 100-Car Field Experiment. Blacksburg, VA, Virginia Tech Transportation Institute: 422.
46. Klauer, S. G., T. A. Dingus, et al. (2006). The Impact of Driver Inattention on Near-Crash/Crash Risk: An Analysis Using the 100-Car Naturalistic Driving Study Data. Blacksburg, VA: 226.
47. Hanowski, R. J., M. Blanco, et al. (2008). The Drowsy Driver Warning System Field Operational Test: Data Collection Methods. Blacksburg, VA, Center for Truck and Bus Safety Virginia Tech Transportation Institute: 295.
48. Olson, R., R. J. Bocanegra, et al. (2009). DRIVER DISTRACTION IN COMMERCIAL VEHICLE OPERATIONS. Blacksburg VA, Center for Truck and Bus Safety Virginia Tech Transportation Institute: 285.
49. Fitch, G. M., S. E. Lee, et al. (2009). Analysis of Lane-Change Crashes and near-Crashes. Blacksburg, VA, Virginia Tech Transportation Institute: 88.
50. Lee, S. E., E. C. B. Olson, et al. (2004). A Comprehensive Examination of Naturalistic Lane-Changes. Blacksburg, VA, Virginia Tech Transportation Institute: 232.
51. McLaughlin, S. B., J. Hankey, et al. (2009). Development of an FCW Algorithm Evaluation Methodology With Evaluation of Three Alert Algorithms. Blacksburg, VA, Virginia Tech Transportation Institute: 120.
52. Knipling, R. R., L. Boyle, et al. (2004). "Commercial truck and bus safety-Synthesis 4- Individual differences and the "high-risk" commercial driver." Transportation Research Board.
53. Kiefer, R. J., D. LeBlanc, et al. (1999). Forward Collision Warning Systems: Development and Validation of Functional Definitions and Evaluation Procedures for Collision Warning/Avoidance Systems.
54. S.J. Brunson, E.M. Kyle, et al. (2002). "Alert algorithm development program NHTSA rear-end collision alert algorithm."
55. Fitch, G. M., H. Rakha, et al. (2008). Safety Benefit Evaluation of a Forward Collision Warning System: Final Report. Blacksburg, VA, Virginia Tech Transportation Institute: 116.
56. McLaughlin, S. B., J. Hankey, et al. (2009). Contributing Factors to Run-Off-Road Crashes and Near-Crashes.

57. Bonabeau, E. (2002). "Agent-based modeling: Methods and techniques for simulating human systems." Proceedings of the National Academy of Sciences of the United States of America **99**.
58. Holland, J. (1988). "The Global Economy as an Adaptive Process." The Economy As a an Evolving Complex System.
59. Artificial Intelligence and Advanced Computing Applications Committee (2007). Artificial Intelligence in Transportation. Transportation Research Board, Washington, DC.
60. Abbeel, P., A. Coates, et al. (2008). "Autonomous Autorotation of an RC Helicopter." 14.
61. Abdulhai, B., R. Pringle, et al. (2003). "Reinforcement learning for True Adaptive Traffic Signal Control." Journal of Transportation Engineering **May/June**: 8.
62. Abdulhai, B., R. Pringle, et al. (2003). "Reinforcement learning for true adaptive traffic signal control." Journal of Transportation Engineering **129**(3): 278-285.
63. Jacob, C. and B. Abdulhai (2006). "Automated adaptive traffic corridor control using reinforcement learning: Approach and case studies." Transportation Research Record(1959): 1-8.
64. Bingham, E. (2001). "Reinforcement learning in neurofuzzy traffic signal control." European Journal of Operational Research **131**(2): 232-241.
65. Choy, M. C., R. L. Cheu, et al. (2003). "Real-Time Coordinated Signal Control Through Use of Agents with Online Reinforcement Learning." Transportation Research Record(1836): 64-75.
66. Adam, Z. M., M. M. Abbas, et al. (2009). "Evaluating Green-Extension Policies with Reinforcement Learning and Markovian Traffic State Estimation." Transportation Research Record(No.2128): 9.
67. Bogers, E. A. I., M. Bierlaire, et al. (2007). "Modeling Learning in Route Choice." Transportation Research Record **2014**: 8.
68. Arentze, T. and H. Timmermans (2003). "Modeling Learning and Adaptation Processes in Activity-travel choice." Transportation **30**: 26.
69. Han, Q. and H. Timmermans (2006). "Interactive Learning in Transportation Networks with Uncertainty, Bounded Rationality, and Strategic Choice Behavior." Transportation Research Board **1964**: 8.
70. Wahba, M. and A. Shalaby (2005). "Multiagent learning-Based Approach to Transit Assignment Problem: A Prototype." Transportation Research Record **1926**: 10.
71. Abbeel, P., A. Coates, et al. (2007). "An Application of Reinforcement Learning to Aerobatic Helicopter Flight." NIPS **19**: 8.
72. Jouffe, L. (1998). "Fuzzy Inference System Learning by Reinforcement Methods." IEEE TRANSACTIONS ON SYSTEMS, MAN, AND CYBERNETICS—PART C: APPLICATIONS AND REVIEWS **28**(3).
73. PTV-AG (2008). "VISSIM 5.10 User Manual."
74. Sutton, R. S. and A. G. Barto (1998). Reinforcement Learning—An Introduction. Cambridge, Massachusetts, USA, MIT Press.
75. Zhang, Y., Xie, Yuanchang, and Y, Zhirui (2007). Development and Evaluation of a Multi-Agent Based Neuro-Fuzzy Arterial Traffic Signal Control System Texas Transportation Institute