A Physics-Based Framework for Estimating Real-Time Platoon Energy Savings

by

Evan Stegner

A dissertation submitted to the Graduate Faculty of
Auburn University
in partial fulfillment of the
requirements for the Degree of
Doctor of Philosophy

Auburn, Alabama
May 4th, 2024

Keywords: truck, platooning, heavy-duty, energy efficiency, vehicle autonomy

Approved by

Mark A. Hoffman, Chair, Assistant Professor of Mechanical Engineering
David M. Bevly, McNair Distinguished Professor of Mechanical Engineering
Scott M. Martin, Assistant Professor of Mechanical Engineering
Roberto C. Molinari, Assistant Professor of Statistics
Brian R. McAuliffe, Senior Research Officer, National Research Council of Canada
Jeffrey LaMondia, University Reader, Professor of Civil Engineering

Abstract

Platooning is a strategy that aims to reduce aerodynamic drag and fuel consumption by driving trucks closely behind each other. However, the advantages of platooning can be compromised by braking loss, which refers to the extra braking needed to follow trucks on hills and in traffic. Conducting experiments to determine the benefits of platooning is time-consuming and resource-intensive, prompting the search for a method to estimate these benefits.

This thesis introduces a framework that allows operators to predict energy savings during platooning, even when braking is involved, providing immediate feedback to platoon operators. The physics-based framework applies to various types of road vehicles. The dissertation outlines practical approaches to implement the framework in real-time, such as a new adaptive estimation of braking loss based on vehicle wheelspeed and a method to query hyperlocal wind data, although the latter did not enhance correlation.

The validity of the framework is confirmed through robust regression analysis of more than 8000 different pairwise comparisons from experimental platoon trials, with the estimated energy change often closely matching the actual energy change and an R^2 value exceeding 0.68 in the model. The proposed methods for estimating braking losses are relatively resistant to errors in drag and rolling resistance, but errors in mass can skew the results by a factor of two. Furthermore, a logistic regression classification approach is introduced to use the framework for making go/no-go decisions, enabling users to specify their desired confidence level. The classification approach demonstrated a 73.5% accuracy when applied to unseen on-road platooning data.

The framework effectively distinguishes the energy impacts of platooning from other vehicle energy consumption. It is grounded in physics, adaptable, capable of real-time operation without the need for a simultaneous reference, sensitive to various parameterizations and signal inaccuracies, and can provide clear feedback to platoon operators on energy savings with a binary outcome. Ultimately, this research can help platoon operators optimize

their energy efficiency by offering realistic fuel savings expectations and guiding decisions about when to engage in platooning.

Acknowledgments

To my advisor, Dr. Mark Hoffman. I don't think either of us imagined we'd be researching semi-truck platoons, but it has been so rewarding. Your guidance, availability, and encouragement taught me more than you know.

To my parents, Chris and Jennifer Stegner. Thank you for your nurturing attitude that encouraged me to be curious and to explore. It seems that childhood desire to be a semi truck driver showed up eventually!

To my friends who were there in the thick of it, rewiring radars and rigging up fuel lines at 2 am, thank you. A special mention goes to Jake Ward. I am so thankful for your uncrushable spirit, shared enthusiasm for rental cars, control wizardry, and above all, your general friendship.

To my daughter Georgia. Your arrival gave me the push to finish writing. I hope you read this someday when you're older, and know that our relationship is worth so much more than any of this work.

To my wife Ashton, for your patience, balance, grace, encouragement, and steadfastness. You saw this through just as much as I did, and I can't count the number of times I had to lean on your help. Look at how much we've grown up! I am so thankful to have you in my life.

Finally, to Eternal God, Jesus Christ. You are the vine, I am the branch. I offer up this research, if for anything, to bring You greater glory. May its failings be mine, and its successes, Yours.

Table of Contents

Abstract ii

Acknowledgments iv

List of Tables ix

List of Figures xi

List of Abbreviations xviii

1 Introduction 1

1.1 Background on Platooning 1

1.2 Research Gap 2

1.3 Contribution 5

1.4 Scope and Limitations 5

1.5 Chapter Outline 7

2 Literature Review 8

2.1 Chapter Overview 8

2.2 Platooning Energy Consumption Research Methods 9

2.3 Internal Factors of Platooning 11

2.3.1 The Effect of Intervehicle Distance (IVD) on Platoons 12

2.3.2 The Effect of Lateral Misalignment on Platooning 16

2.3.3 The Effect of Vehicle Body Shapes on Platooning 17

2.3.4 The Effect of Vehicle Powertrain on Platooning 19

2.3.5 The Effect of Vehicle Speed on Platooning 20

2.3.6 The Effect of Vehicle Mass on Platooning 21

2.4 Platoon Gap Control Strategies 22

2.5 Platooning Disturbances 23

2.5.1 The Effect of Hilly Terrain on a Platooning 24

2.5.2 The Effect of Platoon Interruptions 27

2.5.3 The Effect of the Terrestrial Environment and Other Naturalistic Traffic Wakes . . . 28
2.6 Summary of Various Effects on Platooning Energy Savings . . . 31
2.7 Existing Models of Platooning Energy Savings . . . 32
2.8 Research Gap and Motivation . . . 35
3 Methodology . . . 36
3.1 Platoon Energy Consumption Fundamentals . . . 37
3.1.1 Model Assumptions . . . 39
3.1.2 Choosing the Non-Platooning Reference . . . 42
3.1.3 Assessing the True Platooning Energy Change . . . 45
3.2 Assessing the Adequacy of Available Methods . . . 48
3.2.1 Brief Review of Previous Methods to Infer Benefits . . . 48
3.2.2 Application of Previous Methods to Unseen Data . . . 53
3.2.3 Discussion . . . 56
3.3 Model Formulation . . . 57
3.3.1 Simulation Setup . . . 57
3.3.2 Normalized Power Consumption Calculation . . . 61
3.3.3 Derivation of the Platooning Energy Consumption Function . . . 64
3.3.4 Verifying the Platooning Energy Consumption Function with Simulation 70
3.4 Practical Calculation of Platooning Energy Change . . . 74
3.4.1 Aerodynamic Drag Reduction Calculation . . . 74
3.4.2 Vehicle Acceleration Calculation . . . 79
3.4.3 Braking Power Calculation . . . 83
3.5 Validation Methodology . . . 93
3.5.1 Experimental Datasets . . . 95
3.5.2 Data Processing . . . 105
3.6 Chapter Summary . . . 107

4 Results and Discussion . . . 109
4.1 Sensitivity Analysis . . . 109
4.1.1 Noise Characteristics . . . 110
4.1.2 Sensitivity Analysis Procedure . . . 111
4.1.3 Sensitivity to C_d Error . . . 113
4.1.4 Sensitivity to C_{rr} Error . . . 116
4.1.5 Sensitivity to Mass Error . . . 117
4.1.6 Combined Sensitivity and Summary . . . 119
4.2 Experimental Results . . . 120
4.2.1 Canadian Cut-in Revisited . . . 120
4.2.2 Consolidating Experimental Data . . . 124
4.2.3 Motivation for Robust Regression . . . 129
4.2.4 Representative Result . . . 129
4.2.5 All Results . . . 136
4.2.6 Discussion . . . 143
4.3 A Rudimentary Platoon Energy Savings Classifier . . . 146
4.3.1 Logistic Regression to Provide Estimated Probability of Savings . . . 148
4.4 Chapter Summary . . . 155
5 Conclusions . . . 158
5.1 Contribution . . . 158
5.2 Significance . . . 160
5.3 Recommendations for Future Work . . . 162
5.4 Summary . . . 164
References . . . 165
Appendices . . . 175
A Regression Analysis of the Available Methods . . . 176
A.1 Data Description . . . 176

A.2 Exploratory Data Analysis of Dataset 1 180
A.3 Numeric Variable Selection by LASSO 183
A.4 The Two-Term Model 185
A.5 Intercept Hypothesis Testing 186
A.6 Adding More Terms to the Model 186
A.7 Final Model Selection 188
A.8 Normality of Residuals and Influential Points 189
B The Effects of Net Elevation Change on Platoon Energy Consumption 191
C Simulation Details 195
C.1 Environment 195
C.2 Vehicle Model 197
C.3 Simulation Algorithms 200
D Test Replicates 202

List of Tables

2.1 The effect of trailers on platooning benefits relative to no treatment box trailer, based in part upon the interpretation of [32, 42]. 18

2.2 Original published values for the Schmid et al. drag reduction model [28] compared to the recalculated values [96]. 34

3.1 Interpretation of model coefficients and their 95% confidence intervals. 52

3.2 The parameters used for the simulated ACC controller. 58

3.3 Simulation Mass/Speed/DRR Settings. 59

3.4 Simulation truck and environment parameters. 59

3.5 Reference mean power consumption neural network architecture. 62

3.6 Specifications of the trucks used in the Datasets. 95

3.7 Specifications of Dataset 1 in addition to those in Table 3.6. 98

3.8 Specifications of Dataset 2 in addition to those in Table 3.6. 101

3.9 Specifications of Dataset 3 in addition to those in Table 3.6. 103

3.10 Specifications of Dataset 4 in addition to those in Table 3.6. 104

4.1 The relevant values to calculate $NPC_{inferred}$ and $NFC_{inferred}$ for the Canadian cut-in trial. Values in kilowatts and liters/hour. 121

4.2 The corresponding methods of the "Calculation Method Numbers" in Figure 4.10. 123

4.3 The initial model matrix before data is randomly paired. 127

4.4 The model matrix after data is randomly paired. 127

4.5 The final model matrix for assessing relative platooning energy consumption. . . 128

4.6 The final model matrix for assessing absolute platooning energy consumption. . 128

4.7 The slopes of the planes in Figure 4.13 for inferences on absolute power difference. Slopes which are no different than one are in bold. 132

4.8 The slopes of absolute fuel difference versus inferred fuel difference due to braking and drag differences. Slopes which are no different than one are in bold. The percentages in italics are the percent change of the slope versus its corresponding value in Table 4.7, which is the power equivalent of this table. 134

4.9 The R^2 of each method on $NPC_{true} \sim NPC_{inferred}$. Asterisk indicates the highest R^2 in row or column, and the **bold** values are the average for each 3-by-3 subtable. 137

4.10 The slope of the coefficient related to $\bar{P}_{AD}$ (β_1), expressed as a percentage of the true value. For fuel results, η_{gen} was set to 32.2%. Values in **bold** are those that are not significantly different from 100%, according to the 95% confidence interval. 138

4.11 The slope of the coefficient related to $\bar{P}_{aero}$ (β_2), expressed as a percentage of the true value, with constant air density and no winds. For fuel results, η_{gen} was set to 32.2%. Values in **bold** are those that are not significantly different from 100%, according to the 95% confidence interval. 139

4.12 The slopes of the regression for absolute power and fuel differences, $\Delta P_{true} \sim \Delta P_{inferred}$ and $\Delta F_{true} \sim \kappa \Delta P_{inferred}$. For fuel results, η_{gen} was set to 32.2%. Values in **bold** are those that are not significantly different from 100%, according to the 95% confidence interval. 140

4.13 The slope of each method on $NPC_{true} \sim NPC_{inferred}$ and $NFC_{true} \sim NFC_{inferred}$ (expressed as a percentage). For fuel results, η_{gen} was set to 32.2%. Values in **bold** are those that are not significantly different from 100%, according to the 95% confidence interval. 141

4.14 Six potential logistic regression models for platoon savings classification. 148

4.15 Accuracy of the logistic regression model candidates, averaged across the 18 combinations of DRR, $\bar{P}_{AD}$ adjustment, and weather models. Using probability threshold of 0.5. 150

A.1 Dataset variables. 178

A.2 Summary statistics of the dataset. 181

A.3 Statistics for various candidate models with and without fan power. 188

A.4 Interpretation of model coefficients and their 95% confidence intervals, re-normalized per Equation A.4. 188

C.1 Gear ratios of the simulated truck, to mimic truck A1. 199

List of Figures

2.1 General overview of platooning energy consumption research. 8

2.2 The number of papers referencing different internal factors in [14]. 11

2.3 The mechanism of aerodynamic drag reduction for platooning trucks. 12

2.4 Three-truck platoon fuel savings from [26] with author's insertion of superimposed lead and trail truck benefits. 15

2.5 Fuel savings results from a four-truck heterogeneous platoon [29]. 20

2.6 Aerodynamic drag versus rolling resistance from [44]. 21

2.7 PID versus look-ahead control fuel savings on two sides of a hilly highway. Original graphic using the data in [48]. 24

2.8 Platoons traversing grade are faced with disturbances that are at odds with strict longitudinal spacing. 26

2.9 Marginal effects of various factors on platoon energy savings. Savings are relative to vehicles in isolation with the same factor levels. 31

3.1 Energy use flowchart for vehicles. 37

3.2 The top-level architecture of the proposed platooning energy consumption model. 42

3.3 Normalized Power Consumption versus Normalized Fuel Consumption. Results from simulation as described in Section 3.3.1 and Appendix C. 47

3.4 Figures from [10] demonstrating the previously published method for relating braking to NFC. 49

3.5 Figures from [11] demonstrating the previously published method for relating braking to NFC. 50

3.7 The rolling impact of E_{AD} on fuel savings for the cut-in data of [42], using the approach in [11]. The published true marginal impact of braking is shown for comparison. The intervals shown are the 95% confidence intervals. 55

3.8 NFC_{true} for the cut-in data in [42] versus $NFC_{inferred}$ calculated by [11]. 55

3.9 Discrepancy between the simulation set speed and the actual mean speed. . . . 61

3.10 Brakeless reference mean power consumption results from simulation and the accompanying neural network model to lookup $\bar{P}_{ref}(m, \bar{v})$. 62

3.11 Error contour for the neural network fit of $\bar{P}_{brakeless}(m, \bar{v})$. 63

3.12 Components of velocity over grade. Elevation is not to scale. 65

3.13 The **drag fraction** as a function of mass and mean speed, with no drag reduction (DRR=1.0). Values from simulation. 66

3.14 Mean power consumed for all simulation runs, including the brakeless reference runs. 70

3.15 Normalized Power Consumption versus mass and speed. 71

3.16 NPC_{true} versus $NPC_{inferred}$ for the simulation. $R^2 = 0.998$. 72

3.17 A demonstration of using the OpenWeatherMap API to query wind speed and direction on a 3-hour trip from Auburn to Mobile and estimate vehicle wind speed and yaw angle at $z_0 = 0.5m$ and $z_0 = 1.28m$. 78

3.18 Illustration of the attentuation and phase distortion of the wheelspeed Kalman filter from [103]. 80

3.19 FIR filter errors, showing maximum velocity error in color, and the velocity RMSE as contours. Also shown are the point with the minimum velocity RMSE (green), the minimum acceleration RMSE (yellow), and the selected filter design (red). 83

3.20 Modeled Velocity Drift. 84

3.21 Results of Algorithm 2 with correct parameters. The onset/offset of braking events are marked by gray vertical lines. 90

3.22 Results of Algorithm 2 with incorrect drag and rolling resistance coefficients (Error in $C_d A_f = -5$ m^2, error in $C_{rr} = -0.005$). The onset/offset of braking events are marked by gray vertical lines. 91

3.23 Results of Algorithm 2 with incorrect mass (Error in $m = +10000$ kg). The onset/offset of braking events are marked by gray vertical lines. 92

3.24 Map of the GPS data of Dataset 3. 102

3.25 Auburn CACC System Architecture. 106

4.1 Empirical probability density function of the grade and wheelspeed errors, and the accompanying best-fit T-distribution. Data is derived from all four experimental datasets described in the Methodology. 111

4.2 The impact of the added noise on the simulated wheelspeed, velocity, and grade. 112

4.3 Histograms of the three types of parameter errors, expressed as a percentage of the true value. 113

4.4 The impact of aerodynamic drag error on $NPC_{inferred}$ of a platoon. Parameter errors are uniformly distributed at up to ± 5 m^2, which is 50% of the true value. Percent error is calculated (inferred $-$ true)/true. 114

4.5 The sensitivity of $NPC_{inferred}$ versus true NPC when there is a error in the estimated C_d. $\bar{P}_{AD}$ estimated by RLS method. 115

4.6 The impact of C_{rr} error on $NPC_{inferred}$ and its constituent terms. Parameter errors are uniformly distributed at up to ± 0.005 which is 50% of the true value. Percent error is calculated $\frac{\text{inferred}-\text{true}}{\text{true}}$. 116

4.7 The impact of mass error on $NPC_{inferred}$ and its constituent terms. Mass errors are uniformly distributed at up to ± 5000 kg. Percent error is calculated $\frac{\text{inferred}-\text{true}}{\text{true}}$. 117

4.8 The impact of mass errors on inferred power consumption of a platoon. $\bar{P}_{AD}$ estimated by RLS method. 118

4.9 The net impact of errors on inferred power consumption of a platoon. Inclusive of noise in velocity, grade, and engine power, and parameter offsets up to $\pm 50\%$ for C_d/C_{rr}, ± 5000 kg for mass. $\bar{P}_{AD}$ estimated by RLS. 119

4.10 $NPC_{inferred}$ for the 23m platoons in [42] without cut-ins (top), with cut-ins (middle), and the difference between the two (bottom). Description of the calculation method numbers is in Table 4.2. 122

4.11 The pairwise comparison scheme. 125

4.12 The fitted regression plane of true power difference versus the braking and drag power difference. 131

4.13 The fitted regressions planes for each dataset of absolute power difference versus braking and drag differences. The gridded plane in each panel is the one-to-one plane. The datasets are: Dataset 1: DOE test-track (red, upper left), Dataset 2: NRC (green, upper right), Dataset 3: I-85 (purple, lower left), Dataset 4: Hwy-280 (blue, lower right). 132

4.14 The fitted regression plane of true fuel rate difference versus the inferred fuel difference due to braking and drag reduction. $\kappa = (0.322 \cdot 36e6)^{-1}$ liters per Joule. 134

4.15 A relative platoon energy consumption result, in terms of power (left) and fuel (right). 135

4.16 The slopes of the various methods in absolute terms ΔP and ΔF. The one-to-one line is also shown as a dotted black line. 140

4.17 The slopes of the various methods in relative terms NPC and NFC. The one-to-one line is also shown as a dotted black line. 142

4.18 The confusion matrix for binary platooning classification using $NFC_{inferred} > 1$ as the predictor. 147

4.19 The confusion matrices broken out by truck for binary platooning classification using $NFC_{inferred} > 1$ as the predictor. 147

4.20 The ROC curve for the logistic regression model structure B ($\Pr(NFC_{true} < 1)$ versus $\Delta\bar{P}_{inferred}$) for the training and test set. The true class is equivalent to $NFC_{true} < 1$ in this case. 152

4.21 The three-way confusion matrix for the logistic regression model structure B ($\Pr(NFC_{true} < 1)$ versus $\Delta\bar{P}_{inferred}$) for both the training and test set, with specified acceptable loss rate of 10%. 153

4.22 Trained platoon decision by logistic regression, based on desired two-sided confidence level. The estimated probability of fuel savings is mapped to the left axis and the levels of $NPC_{inferred}$ that correspond to the probabilities are mapped to the right axis. 154

5.1 A demonstration of the kind of information that could be displayed to platoon operators by using this framework. 160

A.1 Calculation process for Normalized Fuel Consumption (NFC), the basis for the target variable NFCr. 177

A.2 A demonstration of the principle behind E_{AD} calculation, showing an instance of braking whereby a truck lost energy. 179

A.3 Pie charts of the data composition. 180

A.4 NFC residual versus various regressors. 182

A.5 Histogram of the NFCr for standalone trucks only. 183

A.6 LASSO regularization to select variables. 184

A.7 Two-term model response plane. 185

A.8 Change in the response due to the addition of fan power. 187

A.9 Plots to check the normality of the least-squares model assumptions. 189

A.10 Supporting figures for the outliers. 190

B.1 Gravity and braking effects on NPC for a fixed drag reduction of 20% and an even power split between drag and other loads. 193

C.1 The elevation and grade profiles of the simulation environment. The elevation and grade profiles from the experimental test track are overlaid. 196

C.2 Screenshot of the vehicle simulation as it passes the 1 km mark and begins the simulated grade. 197

C.3 Screenshot of the simulated trailer loads on the vehicle. 198

C.4 Screenshot of the fuel consumption map settings for the powertrain of the simulated truck. 199

D.1 Dataset 1 test replicates, part 1. 202

D.1 Dataset 1 test replicates, part 2. 203

D.2 Dataset 2 test replicates. 204

D.3 Dataset 3 test replicates. 204

D.4 Dataset 4 test replicates. 205

List of Abbreviations

α	Ratio of reference to platooning η_{gen}
β	Regression coefficient, utilized in various contexts
$\Delta F_{inferred}$	Inferred change in fuel consumption for a microtrip
ΔF_{true}	True difference in vehicle fuel consumption for a microtrip, derived from two separate trials
$\Delta P_{inferred}$	Inferred change in power consumption for a microtrip
ΔP_{true}	True difference in vehicle power consumption for a microtrip, derived from two separate trials
$\eta_{AD,lost}$	Fraction of active deceleration power lost to waste heat
η_{gen}	fuel conversion efficiency
$\hat{a}$	Estimated true acceleration
κ	power-to-fuel conversion factor
drag fraction	Fraction of road load that is spent overcoming aerodynamic drag
ν	Shape parameter of t-distribution, analogous to degrees of freedom
ψ	Effective yaw angle of the vehicle with crosswinds
ρ_a	Ambient air density
ρ_{fuel}	Fuel density
σ	Standard deviation, used in various contexts
a	Instantaneous vehicle acceleration

A_f	Vehicle frontal area
a_{decel}	Fraction of instantaneous vehicle acceleration that is due to active deceleration
a_{model}	Deterministically modeled acceleration without braking
B	Braking boolean used to estimate P_{AD}
C	Average control energy consumption
C_d	Coefficient of aerodynamic drag
C_{rr}	Coefficient of rolling resistance
DRF	Drag Reduction Fraction $(1 - DRR)$
DRR	Drag Reduction Ratio
E_{AD}	Active Deceleration Energy
G	Experimental group for pairwise comparison, divided by truck and track
g	Gravitational acceleration, 9.8 $^{m}/_{s^2}$
LHV	Lower Heating Value
m	Vehicle static mass (excluding linearized rotating mass)
m_{eff}	Vehicle effective mass (including linearized rotating mass)
N_g	In-group index for pairwise comparison
NFC	Normalized Fuel Consumption
$NFC_{inferred}$	Inferred change in the ratio of fuel consumption for a microtrip
NFC_{true}	True change in NFC, derived from two separate trials

$NFCr$	Normalized Fuel Consumption residual ($NFC - 1$)
NPC	Normalized Power Consumption
$NPC_{inferred}$	Inferred change in the ratio of power consumption for a microtrip
NPC_{true}	True change in NPC, derived from two separate trials
P_{accel}	Instantaneous inertial positive acceleration power
$P_{accessory}$	Instantaneous accessory power, including climate control, fan load, and electrical load
P_{AD}	Instantaneous Active Deceleration Power
P_{aero}	Instantaneous aerodynamic drag power
P_{grvt}	Instantaneous gravitational power
P_{plat}	Instantaneous total vehicle power consumption (in platoon configuration)
P_{ref}	Instantaneous total vehicle power consumption (in reference configuration)
P_{rr}	Instantaneous rolling resistance power
PIF	Power Increase Fraction (due to active deceleration)
PRR	Power Reduction Ratio (due to drag reduction)
T	Average test energy consumption
v	Instantaneous vehicle velocity
v_w	Instantaneous vehicle wind velocity in body axis
v_x	Instantaneous vehicle velocity in longitudinal axis
v_z	Instantaneous vehicle velocity in vertical axis

z_0	Surface roughness parameter for wind speed adjustment
A1	Model year 2015 Peterbilt 579 (specifications linked)
A2	Model year 2015 Peterbilt 579 (specifications linked)
T13	Model year 2009 Freightliner M915A5 (specifications linked)
T14	Model year 2009 Freightliner M915A5 (specifications linked)
ACC	Adaptive Cruise Control
ACM	American Center for Mobility
AD	Active deceleration, inclusive of all active methods of slowing a vehicle
API	Application Programming Interface
AUC	Area under the ROC curve
c.ofst	Abbreviation for "Constant Offset" routine
CACC	Cooperative Adaptive Cruise Control
CAN	Controller Area Network
CAV	Connected and Autonomous Vehicles
CDF	Cumulative Distribution Function
CFD	Computational Fluid Dynamics
Constant Offset	Constant offset adjustment routine for P_{AD} estimates
DOE	Department of Energy
DPF	Diesel Particulate Filter
ENU	East-North-Up frame of vehicle coordinates

FHWA	Federal Highway Administration
FIR	Finite Impulse Response, a type of linear-phase digital filter
GPS	Global Positioning System
h.pw	Abbreviation for Hussein power law DRR model
h.rp	Abbreviation for Hussein rational polynomial DRR model
HDV	Heavy-duty vehicles, defined by FHWA as class 7 and 8
IMU	Inertial Measurement Unit
IVD	Intervehicle distance between platoon vehicles
LDV	Light-duty vehicles
MPC	Model-Predictive Control, also known as look-ahead control
NCAT	National Center for Asphalt Technology
NMPC	Nonlinear Model-Predictive Control
NRC	National Research Council of Canada
NREL	National Renewable Energy Laboratory
OWM	OpenWeatherMap, a company providing estimates of hyperlocal weather conditions
PDF	Probability Distribution Function
PID	Proportional-Integral-Derivative
PMG	Automotive testing agency in Blainville, Canada
Q/R	The ratio of process noise to measurement noise in Kalman filtering

RLS	Recursive Least Squares adjustment routine for P_{AD} estimates
RMSE	Root mean squared error
ROC	Receiver Operating Characteristic
ROS	Robotic Operating Software
SAE J1321	Society of Automotive Engineers fuel consumption test procedure
SAE J1939	Society of Automotive Engineers protocol for heavy-duty communications over CAN
sch	Abbreviation for Schmid DRR model
SPN	Suspect Parameter Number, encoded by the SAE J1939 standard
V2V	Vehicle to vehicle

Chapter 1

Introduction

"All models are wrong, but some are useful."

- George Box

1.1 Background on Platooning

By 2020, the United States had almost 3 million registered Class 7 and 8 trucks responsible for transporting more than half of the country's goods [1, 2]. These trucks covered an astonishing 177 billion miles and consumed 28 billion gallons of diesel fuel in 2022 [1]. This underscores the critical role of fuel efficiency in reducing operating expenses in the trucking industry.

Platooning is one of many Connected and Autonomous Vehicle (CAV) technologies that has attracted attention in recent years. Also known as Cooperative Adaptive Cruise Control (CACC) [1], platooning leverages vehicle-to-vehicle (V2V) connectivity, allowing vehicles to follow each other at reduced intervehicle distances (IVDs), thus significantly cutting aerodynamic drag and fuel consumption. In essence, it represents a robotically controlled, collaborative form of "tailgating". Follower vehicles in a platoon can react more quickly and consistently than a typical human driver due to constant IVD feedback and low-latency V2V communication. Platoons typically use the same hardware as Adaptive Cruise Control (ACC) systems, with the addition of V2V communications and GPS. The platoons can be pre-planned or formed opportunistically.

The motivation for platooning may be understood through three overlapping aims:

[1] the term CACC sheds light on platooning's roots in commercial Adaptive Cruise Control (ACC) systems [3], though [4] points out a slight difference between CACC and platooning: platooning was conceived within a "leader-follower" framework, while CACC may be more flexible, with control information for ego vehicles being received from vehicles other than the leader.

1. **Monetary**: The business case for platooning revolves around the actual energy savings of platooning. Reports suggest that platooning has several viable avenues with a payback period less than 1.5 years [5].
2. **Environmental**: Platooning represents another tool in the kit for the reduction of greenhouse gases. A report estimates that platooning could result in a CO_2 reduction of 38 million metric tons over the next five years [6].
3. **Infrastructural**: Platooning has the potential to reduce driver workload and address workforce shortages. In addition, platooning can be integrated into autonomous vehicles, making it a complementary rather than a competing technology.

1.2 Research Gap

Despite the theoretical benefits of platooning, its real-world adoption has been hindered by studies reporting inconsistent energy savings. Theories abound about the unexpectedly low benefits of platooning in real-world scenarios, such as the disruptive effects of other vehicles and road irregularities. These theories remain largely untested and unquantified.

Theories involving platoon disturbances and traffic irregularities have not been tested largely due to their real-world unpredictability. The uncertain nature of practical road conditions makes it difficult to accurately assess the benefits of platooning using only traditional energy consumption metrics. The industry standard is a dedicated test campaign (such as SAE J1321 and J1526), involving at least three trucks, multiple replicates, hours of warm-up, and gravimetric fuel weighing, among other things. These dedicated test campaigns are required to obtain a high level of statistical confidence in fuel savings, but are very time-consuming and resource-intensive.

This highlights an opportunity. Existing methods for platooning benefit determination cannot operate in real time. If they could, it would remove some barriers to platooning adoption. Two potential use cases of in-situ benefit estimation come to mind, the first being gamification and the second being optimization.

The potential for gamification in platooning is significant, particularly for opportunistic platooning, where platoons form spontaneously[2]. Gamification refers to the addition of game-like elements in a non-game context to encourage participation in a feature. Gamification is often implemented in eco-driving strategies to encourage their use[3][8]. Platooning, as with any other human-initiated feature, requires a compelling reason for its utilization. However, platooning faces especially steep obstacles to user participation:

- **Competing Technology:** From an operator's standpoint, platooning is not apparently better than Adaptive Cruise Control (ACC). In fact, in some ways, ACC is more appealing: Platooning at small IVDs requires a surrender of control and visibility that can be unappealing [9].
- **Required Dual Participation:** In contrast to many eco-driving technologies, platooning requires the participation of multiple operators. To create various platoons on the road, drivers would need to modify their driving habits by accelerating or decelerating to synchronize with nearby vehicles. It seems unlikely that a driver would proactively seek to form a platoon with another driver if they did not perceive any concrete advantages from doing so. Even if an individual operator desired platooning benefits, they would still have to find another driver willing to be in the platoon. The leader has even less incentive to join the platoon, as the benefits are typically smaller.

The platoon payback period is strongly dependent on driver participation [5]. If both drivers received useful feedback on the savings from platooning, it would increase their desire to create a platoon and their acceptance of the technology. Therefore, the implementation of a feedback system would add a tangible benefit to platooning and could significantly improve platooning adoption.

In addition to gamification, optimization represents another important use of accurate in-situ platooning energy predictions. The importance of optimizing the energy usage of

[2]For a more thorough explanation of opportunistic platooning, see [7].

[3]Some examples include dashboard lights that illuminate with energy-efficient throttle input, or detailed energy flow graphics displayed in some hybrid vehicles.

Heavy Duty Vehicles (HDVs) has been well recognized. Typically, this energy optimization is carried out at the final metric level. By having a prediction of the energy advantages of platooning, a more detailed, sophisticated, and understandable platooning energy management can be achieved. Accurate energy optimization is even more important for electric vehicles. Electric HDVs are especially range- and time-constrained due to payload limitations and long charging times. Thus, electric HDVs present an excellent target for optimization. With an accurately characterized impact of platooning, such vehicles could leverage not only energy savings but also significant time savings from reduced charging. The native connectivity and software-defined characteristics of many electric vehicles make them ideal for the integration of platooning technology.

Current platoon energy consumption models are not structured for the kind of real-time feedback that could gamify or optimize platooning in complex real-world situations. This is because existing platoon energy consumption models are almost entirely focused on the drag reduction aspect of platooning and have not incorporated the impact of braking losses.

The benefits of platooning hinge not only on drag reduction, but also on braking losses, which are more accurately termed "active deceleration (AD) losses" (to encompass all forms of actuator-induced deceleration, including regenerative braking, engine retarders, and friction brakes) [10, 11]. In this dissertation, the term "braking losses" will be used interchangeably with the more precise technical term "AD losses"[4].

In the literature review, it will be shown that braking losses have a significant impact on platooning energy consumption, and crucially, that there is currently no way to account for the impact of these braking losses on the actual energy savings of platooning. This means that platooning operators have no way of knowing whether platooning is saving or losing fuel.

[4]Nonetheless, in mathematical expressions, the subscript "AD" will consistently be employed to specify that all types of active deceleration are encompassed in the braking power assessments.

1.3 Contribution

In light of the research gap, this dissertation will present a framework that enables platoon operators to infer energy savings during platooning engagement even when there is braking. This platooning energy inference seeks not to measure but to infer the benefits of platooning, a distinction that highlights the practical value of the research.

To fulfill its purpose, the following characteristics must be demonstrated by the framework:

1. **Energy Impact Isolation:** The framework must clearly distinguish the energy effects of platooning from other vehicle energy uses, ensuring the accuracy of the inference.
2. **Adaptable:** The framework must be easily adapted to future vehicle architectures, otherwise it would quickly become obsolete.
3. **Real-Time:** The framework must function in real time without a concurrent baseline reference.
4. **Known Sensitivity:** The framework's sensitivity to errors and disturbances must be characterized, particularly parameter errors such as mass and rolling resistance.
5. **Feedback Mechanism:** The framework must provide useful "go/no-go" feedback.

The "framework", as we shall call it, should be demonstrated by the end of this dissertation to meet all of the above criteria. It will be modular and physics-based, and it will be verified using both simulated and experimental data. The objective is not to reevaluate the benefits of platooning from previous experimental studies, but to demonstrate that the conclusions regarding energy consumption from the framework are valid.

1.4 Scope and Limitations

Thus, the meaning of the title of this dissertation, "A Physics-Based Framework for Estimating Real-Time Platoon Energy Savings", begins to become clear. The framework is "physics-based", as ultimately the inferences that will be drawn by the model are rooted in

vehicle energy demands in a Newtonian sense, as opposed to empiricism. It is a "framework" because it is modular, allowing the submodels within to be exchanged appropriately. It is "real-time" in that it can provide feedback during operation. The concept of "real-time" is in contrast to "post-processing" more than it is to "a slight delay", though for optimization purposes we would like any delay to be as small as possible. For "platoon energy savings", "energy" is taken to mean either energy consumed to move the vehicle (at the tires) or energy consumed at the prime-mover level of the powertrain (be it fuel or battery). Meanwhile, "platoon savings" refers to the isolated effect of platooning versus not platooning, with all other factors being equal.

The inferred platoon savings provided by this framework will not claim to demonstrate a level of precision comparable to that of a regimented fuel test (SAE J1321). We need not go that far; it is enough to demonstrate that the inferences are correlated with actual savings, using simulation as a basis for sensitivity analysis and experimental results for verification.

This dissertation will not cover all the extensive research questions related to platooning. For example, further issues related to platoon planning will not be discussed, despite the role that the proposed framework would play in opportunistic platoon planning. Instead, this study concentrates on a more specific issue: the use of energy during a "platooning microtrip", which is defined as a segment of a trip in which platooning is actively engaged.

In this dissertation, the focus is on heavy duty vehicles (HDVs) as opposed to light duty vehicles (LDVs). Although light-duty vehicles (LDVs) can also save energy through platooning, most of the research is focused on HDVs for a couple of reasons:

- **Fuel Consumption:** HDVs consume much more fuel per vehicle mile traveled and travel many more miles per vehicle. The benefit of platooning two trucks is much more substantial than the benefit of platooning two cars.
- **Duty Cycle:** Line-haul HDVs spend much more of their duty cycle on interstates and highways, where platooning is most beneficial.

Although the majority of platooning research studies are centered on HDVs [12], the findings can be extrapolated to LDVs, albeit on a different scale. With this in mind and the stated goal for adaptability, the framework proposed in this dissertation can easily accommodate LDVs.

HDVs can vary in shape and size, particularly due to the interchangeability of trailers. The data used in this study is limited to trucks hauling box trailers; therefore, the findings will be specifically applicable to trucks with box trailers. However, the framework should naturally accommodate other types of trailer by adjusting the aerodynamic drag reduction models.

1.5 Chapter Outline

Having established the background, motivating factors, and objectives of the framework, this dissertation is structured as follows:

- **Literature Review:** To provide a more thorough understanding of platooning energy consumption, and factors affecting it, the current research understanding of platooning vehicles will be explored.
- **Methodology:** The theory, development, practical implementation, and validation process of the framework will be reviewed.
- **Results and Discussion:** The sensitivity of the framework will be investigated through simulation, and experimental data from various origins will be integrated to verify the framework. The implications of the findings will be discussed and a basic classifier will be presented.
- **Conclusions:** It will be assessed whether the framework satisfies the requirements that have been laid out in this Introduction and a summary and suggestions for future research will be given.

Chapter 2

Literature Review

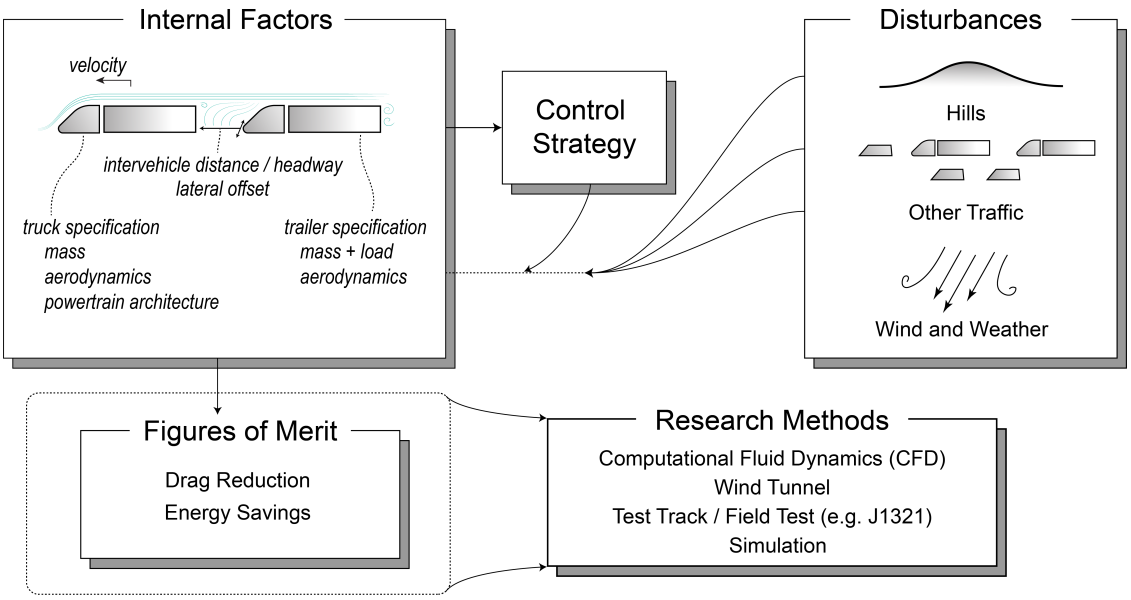

Figure 2.1: General overview of platooning energy consumption research.

Before a framework to infer platooning benefits is derived, it is helpful to understand fully what has been done to research platooning. Platooning has been explored as a way to save energy, both in terms of individual vehicle benefits and in terms of benefits to the broader transportation system. This dissertation is primarily concerned with determining the benefits of individual vehicles, so that will be the focus of this literature review.

2.1 Chapter Overview

Figure 2.1 illustrates the various research challenges and approaches in regard to platooning energy consumption, and also serves to demonstrate the structure of the chapter.

First, a quick overview of the methods and figures of merit for platooning research will be given. The chapter will then begin by exploring the effect of "Internal Factors" on platooning energy consumption. Internal factors are defined as characteristics that make up the platoon, such as the type of trucks used, the setpoint of the intervehicle distance (IVD), and so on.

Once a good understanding of internal factors has been established, the focus will shift to how disturbances affect platoon energy consumption. Disturbances discussed include hilly terrain, platoon interruptions, and the effects of wind and weather. The impact of disturbances on platoon energy consumption is tightly coupled with the platoon IVD control strategy, so that will also be discussed.

After that, the chapter will introduce existing models for platooning fuel savings, which are primarily aerodynamically based. This literature review will show that the impact of disturbances on platooning energy savings is not sufficiently addressed in the literature. Specifically, no platoon energy consumption model includes the absolute impact of controller-induced dynamics, and the understanding of how platoon intrusions affect the platoon savings is limited.

2.2 Platooning Energy Consumption Research Methods

Broadly speaking, prior investigations of platooning energy savings adopt one of three approaches:

1. **Aerodynamic-focused:** The aerodynamics of platooning are directly invesigated, through Computational Fluid Dynamics (CFD), wind tunnels, and in-situ measurements.
2. **Energy-focused:** Test track and field tests are conducted to ascertain the resultant energy benefits of platooning, typically fuel.
3. **Simulations:** Simulated platoons are subjected to disturbances and their responses analyzed, commonly with the goal of improving the robustness of the platoon to dynamic disturbances.

Each method has its proper place. CFD studies allow very fine insights into the nature of aerodynamic flow in and around platooning vehicles, but it can be difficult to accurately model these features, and simulations can be very resource intensive. Wind tunnel testing

requires a dedicated experimental setup but has been used very successfully to investigate the aerodynamics of various types of platoons. Field and test track studies provide insight into the practical platooning benefit, but are subject to many uncontrollable variables. Additionally, they require many man-hours and replicates and thus may be seen as the most expensive of the methods for platoon testing. The recommended methodology for fuel testing on the test track is a modified SAE J1321 test [13], which was originally intended for testing modification of a single vehicle. Finally, simulations are dependent on accurate vehicle modeling (and therefore the other methods) to properly represent the platooning energy consumption but can provide insight about disturbances much more naturally than the other methods. The representation of aerodynamic drag reduction in simulations is itself dependent on data from CFD, wind tunnels, and field measurements. All of these methods have been employed to determine how vehicle platoons consume energy. Let us first look at what has been gleaned about the effect of internal factors on platooning energy consumption.

2.3 Internal Factors of Platooning

Internal factors of platooning refer to the characteristics inherent to a platoon. For example, truck type is considered an internal factor, whereas external factors such as wind strength and direction are not included. The distribution of studies per internal factor is shown in Figure 2.2, adapted from [14]. The graph illustrates that a significant portion of platooning research has focused on analyzing how the separation distance affects energy savings in platooning. Therefore, we will start with IVD first, then proceed to the other internal factors.

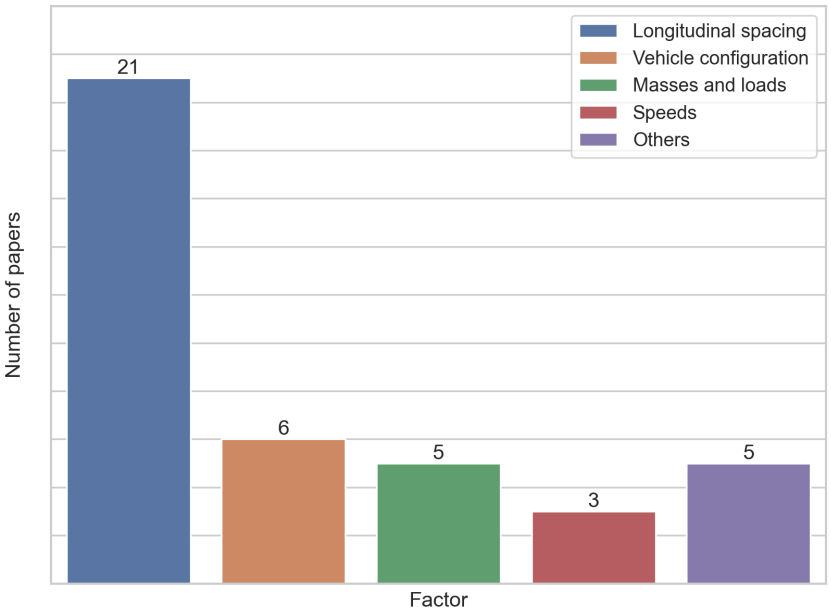

Figure 2.2: The number of papers referencing different internal factors in [14].

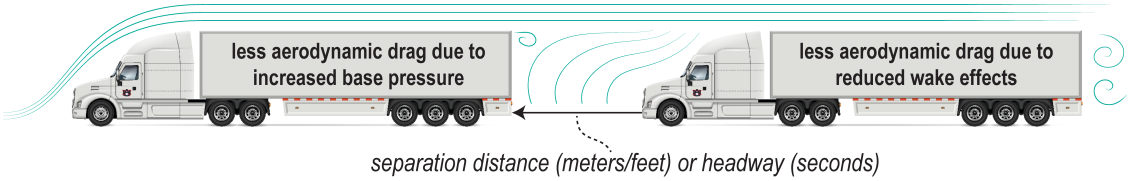

Figure 2.3: The mechanism of aerodynamic drag reduction for platooning trucks.

2.3.1 The Effect of Intervehicle Distance (IVD) on Platoons

The first internal factor that will be investigated is IVD. IVD is just one of several ways to describe the vehicle-to-vehicle distance of the platoon; so, some definitions are in order. The IVD of platoons may also be referred to in the literature as longitudinal spacing, headway, separation distance, or time gap. IVD may be classified spatially or temporally. In the case of constant speed platooning, time gaps and spatial gaps are equivalent, although once vehicles begin to change speed, additional definitions are required [15]. Henceforth, IVD is taken to mean the actual spatial distance between vehicles, unless further distinction becomes necessary.

As Figure 2.3 shows, the mechanism of drag reduction is quite different for leading and trailing vehicles. For lead vehicles, the stagnation area on the front of the trailing vehicle increases the base pressure on the back of the leader's trailer, effectively pushing it [16]. As a result, drag savings increase monotonically as the IVD closes[1]. The increased base pressure effect occurs in a limited range compared to the wake shielding experienced by trailing trucks. Generally speaking, the effects of increased base pressure for lead vehicles are mostly absent for headways larger than one vehicle length (75 feet / 22.9 m) [9, 17].

Aerodynamic advantages for following vehicles can be experienced over distances significantly greater than the length of a single vehicle and may extend beyond 320 feet (100m) [9]. However, the reduction mechanism is more complicated than the increased base pressure effect. For trailing trucks in zero-yaw, aligned conditions, the following aerodynamic mechanisms occur in order of decreasing importance [16]:

[1]There is a limit to the trend, as the truck gets exceedingly close (0.5 meters), a very strong vortex can form and increase the drag of the lead truck [16]

1. Reduced overall oncoming wind velocity, termed "wake shielding" (reduces drag)
2. Decreased accelerations around the frontal edges (increases drag)
3. Flow impingement on the front radii at close follow distances (increases drag)
4. Reduced roof air deflector efficiency (increases drag)

Overall, wake shielding dominates the other effects of IVD on following vehicle drag reduction, and generally, as the separation distance decreases, the energy savings of trailing trucks increase. However, some interesting and counterintuitive aerodynamic behavior occurs for trailing vehicles (and middle vehicles) once the separation distance is less than one vehicle length. Namely, the following vehicle experiences reduced aerodynamic benefits. For HDVs, this reduction in benefit typically occurs below about 40 meters. This reduction occurs in both wind tunnel drag measurements and fuel savings measurements for trailing vehicles [18]. The reduction in benefit was initially hypothesized to be due to several phenomena:

- **Aerodynamics:** An unfavorable aerodynamic regime occurs for followers in the 30' to 50' range, also known as flowfield resonances [19]. Flow impingement, decreased accelerations, and increased drag of the truck-trailer gap appear to be the culprit, at least for cabover-style trucks [20].
- **Driver lateral behavior:** At very close following distances, drivers of manually-steered vehicles may become uncomfortable with the lack of visibility, and move out of alignment with the lead vehicle. Lateral offset effects will be discussed in further detail later in this chapter.
- **A reduction in available cooling air:** Very close following distances may not allow sufficient airflow to the radiator, forcing the engine fan to come on. In [21], it was postulated that platooning may reduce engine cooling airflow to the point that an

engine fan must turn on. The engine fan is quite a large power draw for heavy-duty trucks (15-30 kW per [22]), and could compromise the energy benefit of platooning. [2]

Recent literature has mostly settled on the causality of trailing truck energy benefit reduction. The true nature of the benefit reduction is rooted in aerodynamic drag, although lateral offset and thermal considerations may further reduce the benefit in practice.

Generalization to platoons of more than two vehicles

The aerodynamic savings for two-vehicle platoons easily extend to longer platoons. It turns out that the lead vehicle of a three-vehicle platoon is not affected by the presence of the third vehicle and the third vehicle is relatively unaffected by the position of the lead vehicle [25, 17].

Vehicles in the middle of the platoon experience a superposition of leading and trailing vehicle aerodynamic effects. A mid-platoon vehicle has both wake-shielding from the vehicle(s) preceding and increased base pressure from the vehicle behind it. This superposition of leader and follower benefits can be seen in many of the prior art dealing with 3+ truck platoons, but it is exemplified in the fuel savings results for three-truck platoons.

[2]Indeed, in a 2015 CFD study [23] it was found that at a 5-meter gap, cooling airflow is more than halved for follower trucks, and the engine fan was likely to experience more uptime. In [24] the reduced power demands while platooning were incorporated into a CFD simulation of the reduced cooling airflow of a platoon follower.

To demonstrate the superposition of lead and trailing truck benefits on middle trucks, Figure 2.4 shows the fuel savings of three-truck platoons of [26]. In Figure 2.4, the original middle truck fuel savings have been overlaid with the author's superposition of the lead and trailing truck benefits following the form $\Delta F_{middle} = \Delta F_{Lead} + 0.7\Delta F_{Trailing}$, which can be interpreted as 100% of the lead truck fuel savings and 70% of the trailing truck fuel savings.

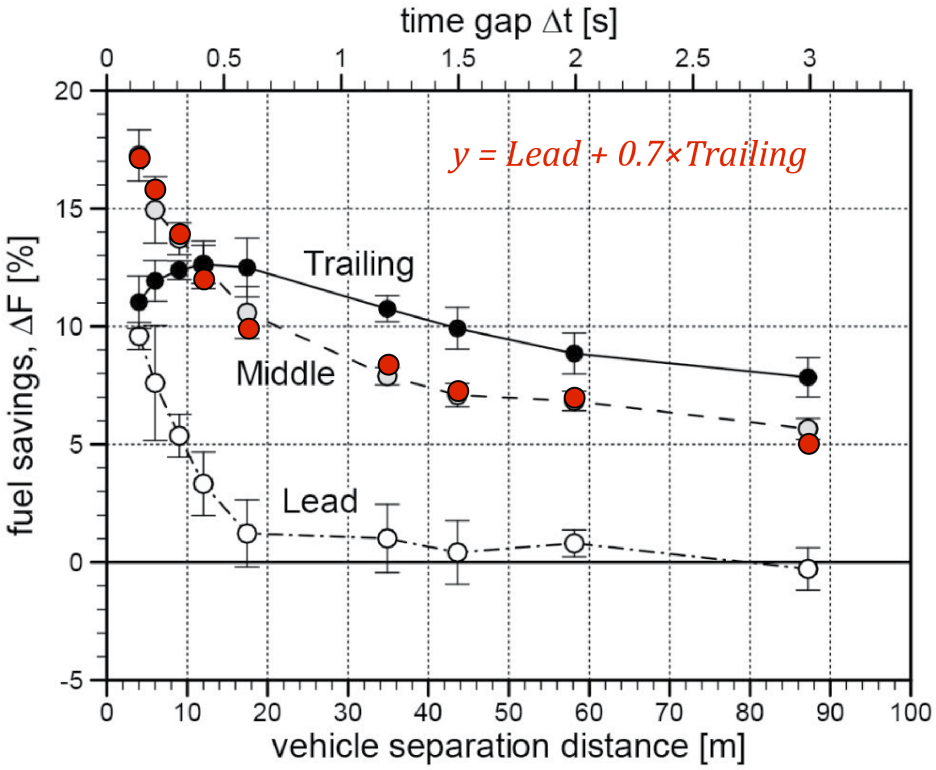

Figure 2.4: Three-truck platoon fuel savings from [26] with author's insertion of superimposed lead and trail truck benefits.

Beyond a three-vehicle platoon, each subsequent following truck has an increased wake deficit [27, 28], up to a limit. The author's work characterized the fuel savings of a four-truck platoon experimentally, finding that the savings were similar in nature to those of a three-truck platoon [29].

This concludes the discussion of the impact of IVD on the energy consumption of platooning. It was seen that IVD has a very strong influence on the magnitude of energy savings

of a platoon, and that this effect is fundamentally different for leaders and followers. Next, the impact of lateral misalignment on platoon energy savings will be explored.

2.3.2 The Effect of Lateral Misalignment on Platooning

Just as platoons are sensitive to the size of the space from truck to truck, they are also sensitive to the alignment of trucks, or conversely the lateral misalignment or offset. It is more difficult to generalize how a lateral misalignment affects the platoon's aerodynamic performance. Within this subsection, the impact of lateral misalignment will be investigated from a zero-mean wind yaw condition.

As with IVD, the aerodynamic behavior for lateral misalignment is quite different for leading and trailing vehicles. The leading vehicles are insensitive to lateral misalignment until very small IVDs. Lateral offset affects trailing trucks in a much more complex manner and will be revisited in the external factors section on wind and turbulence [30].

Small in-lane lateral adjustments from everyday driving are unlikely to impact platooning energy savings measurably [31]. In wind tunnel research, lane offsets of up to 1.3 meters had marginal impacts on the drag reduction of a platoon [32, 17].

Track-based fuel testing of lateral misalignment was conducted in 2019 on the PMG test track[3]. The vehicles were driven in three conditions: aligned (no offset), half-offset (0.65-meter offset), and full offset (1.3-meter offset).

The offset was manually maintained using the centimeter-level accurate GPS position solution displayed to the driver of the following truck, resulting in offset errors ± 0.2 m. Data from these trials was investigated from a fuel savings perspective, and later from an aerodynamic perspective [31, 33]. From these results, the maximum lateral misalignment was found to measurably decrease absolute fuel savings by up to 3-4%.

The cooling air supply is much more sensitive to lateral misalignment than drag savings are, making intentional lateral misalignment an easy way to alleviate platooning thermal

[3]This is the same campaign which will be used in the results to form Dataset 2, although these offset tests were not available for analysis.

concerns at short spacings [34]. Intentional lateral misalignment has also been suggested to minimize rutting of the roadway by platooning vehicles and prolong the life of the pavement [35, 36], although it has been suggested that this is not something to be concerned about for gaps larger than 0.2 s [37].

In general, the consensus on lateral misalignment is that small offsets minimally reduce the benefit of platooning. Additionally, wake shielding effects display complex behavior versus lateral misalignment and are more sensitive to lateral offset than base pressure effects. Having discussed the influence of lateral offset and IVD, we now shift our focus to how the shapes of vehicle bodies affect the energy consumption of platooning.

2.3.3 The Effect of Vehicle Body Shapes on Platooning

Platoons in the real world, especially those that form opportunistically, are likely to be heterogeneous, with varied body and trailer styles. Many researchers have explored how various aerodynamic devices and body shapes impact the platooning benefit. HDVs are a combination of truck and trailer, however, each of which takes many forms. Although trucks and trailers have a combined influence on aerodynamics, their influence will be explored individually as far as possible.

Tractor Body Styles

In North American studies, the majority of studies have investigated engine-forward tractors with sleepers. In contrast, European and Asian platooning research has used cab-over-engine tractors. To the author's knowledge, little work has been done to harmonize the results from North American tractors with those from European/Asian tractors. Cab-over trucks show greater potential energy savings due to their shorter length, which places the leading and trailing trailers closer together for the same spacing [38]. When tractor body shapes are mismatched, it has been suggested that the higher savings are found with the

more aerodynamic vehicle in front, for both two-truck platoons [39] and four-truck platoons [40].

Trailer Styles

The impact of trailer style on a platoon has been relatively well studied. Truck trailers take many forms, however, from box trailers (dry vans) to flatbeds to tankers. The box trailer (or dry van) is the most common type of truck trailer, so naturally it has received the most research attention. Box trailers may have several different types of aerodynamic treatment, including side skirts and boat tails. Of the two, the differential platooning benefits are greater with boat tails than skirts [41]. It appears that the combination of boat tails and trailer skirts preserves platooning energy savings [14], although some results are normalized to a vehicle without aerodynamic devices, which mixes the device benefits with the platooning benefits. The presence of skirts alone on trailers reduces the potential benefits of platooning [32].

The overall impact of boat tails and skirts has been summarized in Table 2.1. As is typically the case, leader vehicle energy savings are less sensitive to trailer aerodynamic treatments than follower savings are [32].

Table 2.1: The effect of trailers on platooning benefits relative to no treatment box trailer, based in part upon the interpretation of [32, 42].

Aerodynamic trailer treatment	Effect on potential platooning benefits
Plain box (no skirts or tails)	baseline
Skirts	reduces benefits
Boat tails	maintains/increases benefits
Skirt + boat tails	complex effect, further study needed
Lead flatbed	reduce lead/increase follow benefits
Follow flatbed	similar benefits

The influence of flatbed trailers on platooning fuel savings has also been investigated. Placing a flatbed trailer on the trailing truck has a minimal impact on the platoon fuel savings relative to the platoon benefit with box trailers, but a flatbed in the lead reduces

the benefits of the lead truck and potentially increases the benefits of the follower. Having briefly explored the impact of various body styles and trailers on platoon savings, we turn to the vehicle powertrain.

2.3.4 The Effect of Vehicle Powertrain on Platooning

Another internal factor in platooning is the architecture of the vehicle's powertrain. For our purposes, the vehicle powertrain can be thought of as a black box that produces mechanical power from one or more energy sources. The most common energy sources for vehicle powertrains are fuel and electrical energy stored in batteries.

The energy input of a vehicle powertrain is not exactly linear with its power output. One consequence of this fact is that aerodynamic drag reduction is not precisely linear with fuel savings. However, it is still common to assume a linear change in fuel consumption with reduction in aerodynamic drag [43], and the clear correlation between wind tunnel savings and savings on test track attests to the validity of the assumption.

With respect to platooning, vehicle powertrains have the same duty to produce mechanical energy, but power demands are reduced by the aerodynamic drag reduction of platooning. There are some important cases when the powertrains non-linearly impact the platoon benefits:

- When a vehicle must actively decelerate due to controller-induced dynamics, some powertrains can reclaim a portion of this energy through regenerative braking, whereas friction/engine braking discards that energy entirely.
- When radiator airflow is reduced by platooning, sometimes the fan duty cycle must be increased to meet cooling demands. HDV engine fans consume a large amount of power, up to 40 horsepower [22].
- When powertrain output is saturated, platooning may lead to a different operating condition. If platooning enables or forces the selection of a different transmission gear, the efficiency of both engine/transmission will change. One hypothetical case where

this may occur is on an uphill segment, where platooning reduces the total road load just enough to allow a lower gear to be run.

If the vehicles in the platoon have heterogeneous powertrains, then the highest-consuming powertrains will unsurprisingly have a greater impact on the platoon team benefits. In Figure 2.5 of [29], the second and third trucks were much less fuel efficient than the first and fourth trucks. Ultimately, fuel savings were dominated by the second and third trucks.

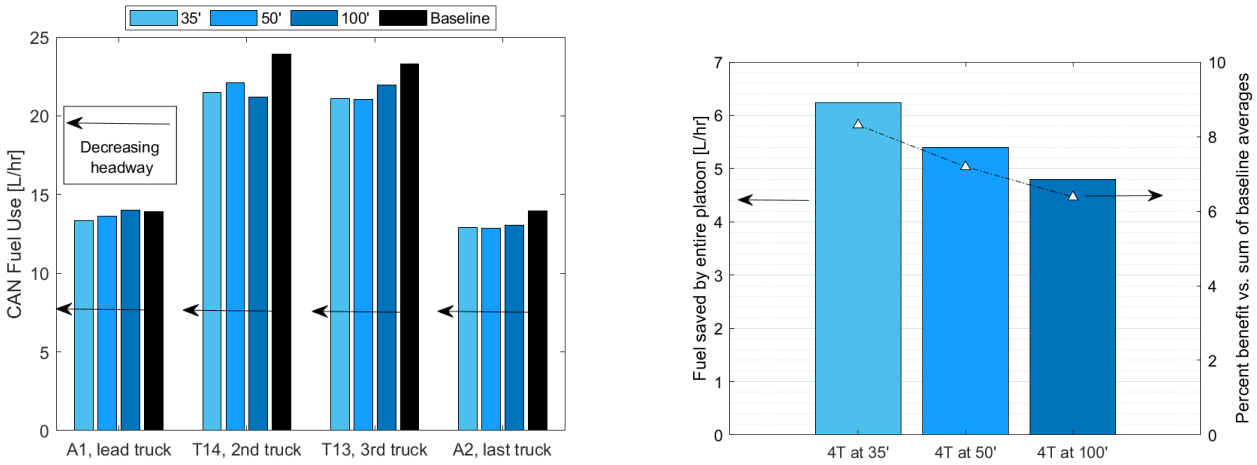

(a) Absolute fuel consumption.

(b) Absolute and relative fuel savings.

Figure 2.5: Fuel savings results from a four-truck heterogeneous platoon [29].

In summary, in many cases, the impact of non-linear vehicle powertrains may be neglected, but thermal considerations, saturation, and speed variations may cause the required power output to be nonlinear with aerodynamic savings. Next, the impact of platoon speed on energy savings will be covered.

2.3.5 The Effect of Vehicle Speed on Platooning

As speed increases, the aerodynamic drag of the vehicles increases quadratically and other road loads increase more slowly. As Figure 2.6 shows, vehicle parameters determine when aerodynamic drag increases in comparison to other road loads [44].

A 64 mph (28.6 m/s) platoon saved 42% more fuel than a 55 mph (24.5 m/s) platoon in one simulation [45]. Thus, as speed increases and drag becomes more dominant, potential

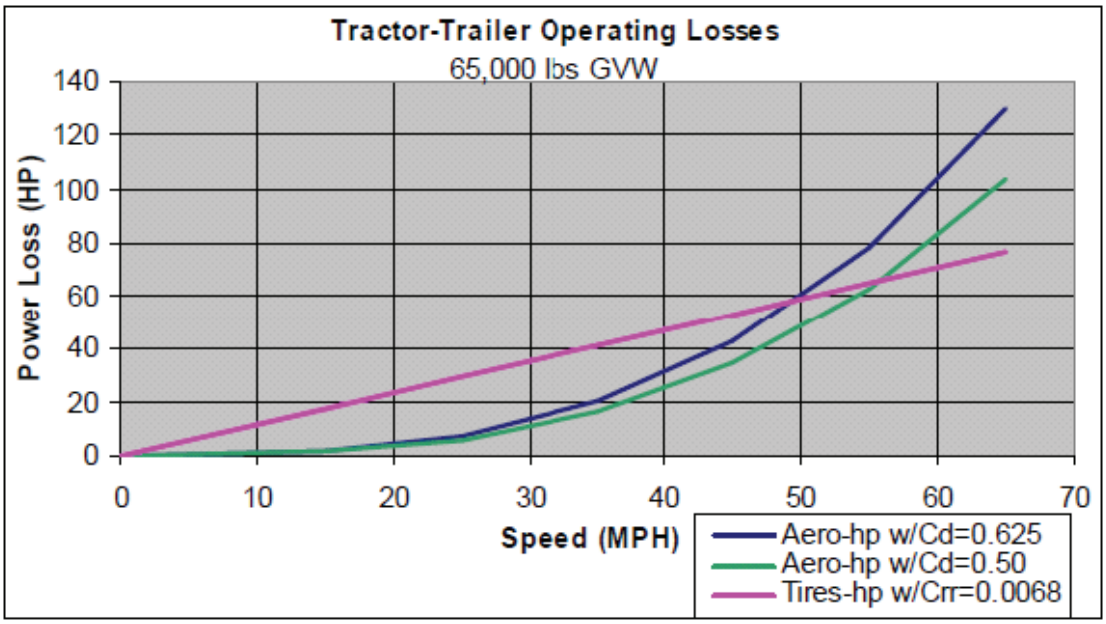

Figure 2.6: Aerodynamic drag versus rolling resistance from [44].

platooning benefits do as well. The expectation that platoon savings will increase with increased speed has not been experimentally confirmed in the literature [21], and it has even been suggested that the percent savings do not measurably change as a function of speed [46].

2.3.6 The Effect of Vehicle Mass on Platooning

Increases in vehicle mass have the opposite effect as increases in speed: increased mass will decrease the potential for fuel savings. This is rooted in the fact that the mass of the vehicle is proportional to rolling resistance. If additional speed "concentrates" platoon energy savings, the additional mass "dilutes" the energy savings [47].

However, this is far from the whole story. Mass increases dilute savings for constant-speed platoons, but when speed variations and especially grade are introduced to the platoon, vehicle mass has a much more complex effect on the energy savings of the platoon. The relationships governing the effect of vehicle mass on energy savings will be explored in more detail in the Disturbances section (Section **??** which we will turn attention to shortly.

However, before that can be done, the various platooning gap control strategies must be overviewed.

2.4 Platoon Gap Control Strategies

Ontologically, gap control is an internal factor of platooning, but the effect of the gap control is only related to disturbances. The effect of the platoon control strategy on energy savings is embedded in the platoon disturbances. In fact, in the case of constant speed, level ground platooning, or even over gentle terrain, it does not matter what type of platoon control is being used for energy benefits [48]. All significant differences in the energy efficiency of gap control strategies come from the handling of disturbances.

The simplest type of platoon gap control is classical linear control. This type of platoon control linearly controls the IVD of the platoon, most often using a proportional integral derivative (PID) control. Linear gap control seeks to accurately track a reference longitudinal spacing. Many methods have been proposed to guarantee that a platoon of N vehicles will not have increasing acceleration demands for each subsequent vehicle, a concept which has been formalized as "string-stability" [49]. Classical linear control is powerful and flexible for its simplicity, capable of handling a wide array of disturbances. However, classical control is not without limitations. In particular, classical control cannot anticipate situations that are energy-suboptimal. For this, researchers have turned to look-ahead control, also known as Model Predictive Control (MPC).

Look-ahead control for platooning naturally arose out of eco-driving methods for single vehicles. Eco-driving, also known as predictive cruise control, uses a preview of the upcoming roadway to optimize the dynamics of the vehicle. Eco-driving is all-encompassing of both roadway and traffic information. One main approach to eco-driving is to optimize the speed profile, as in [50].

Look-ahead control for platooning includes the additional goal of following a close distance to gain platoon benefits. There are numerous approaches to this. Some approaches

coordinate the speed profiles for the entire platoon (a "platoon-layer" coordinator) [51, 15], while others coordinate each vehicle's profile separately [48, 52]. To the author's knowledge, the vast majority of IVD gap control is accomplished through either some form of lookahead control or classical control. The exact details of how these different control strategies impact fuel consumption will be the focus of the next sections.

2.5 Platooning Disturbances

In this section, we explore the impact of different disturbances on the energy of platooning, or at least the existing knowledge on the subject. These disturbances encompass hilly landscapes, interruptions in traffic flow, and aerodynamic forces caused by other vehicles and terrestrial winds. The discussion will proceed in the mentioned order, starting with an overview of platoon control methods due to their relevance to speed and gap regulation within the platoon. Subsequently, we will delve into the latest research on platooning in hilly terrains. The section will then investigate the impact of traffic interruptions on platooning dynamics. Lastly, we will analyze how environmental conditions, such as air composition, mixed traffic scenarios, and wind-induced yaw effects, influence the aerodynamics of platooning.

2.5.1 The Effect of Hilly Terrain on a Platooning

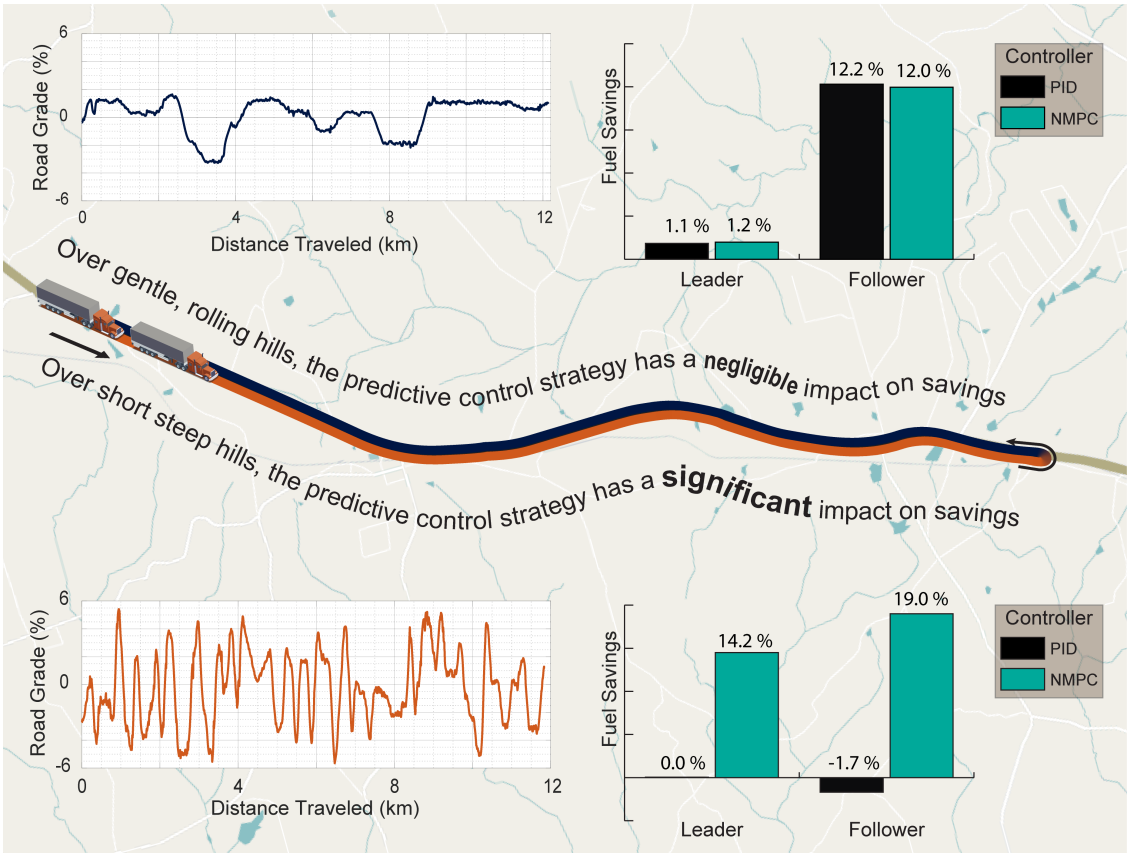

Figure 2.7: PID versus look-ahead control fuel savings on two sides of a hilly highway. Original graphic using the data in [48].

Hilly terrain is a very important factor to consider for platooning. As Figure 2.7 shows, the impact of hills is complex and depends on both the physical topography of the hills as well as the IVD gap control strategy. Several on-road trials of platooning systems have resulted in unexpectedly low fuel savings as a result of grade [53, 54, 55, 56]. Platooning benefits are more sensitive to road grade than isolated vehicle fuel consumption is, and a thorough understanding of this sensitivity is required to make any claims on a platoon's efficiency on real roadways.

Provided that hills are not steep, the most energy-efficient driving strategy is constant speed [4][58]. A **steep** hill is one that either ascends so quickly that it exceeds vehicle power

[4]There are some strategies such as pulse-and-glide [57], which can achieve better efficiency for some duty cycles.

limitations, or the gradient descends rapidly enough to require braking. Steepness is vehicle dependent, and will not usually be symmetric for uphills and downhills. This asymmetry is because steep downhills are determined by rolling resistance, mass, and aerodynamic drag, but steep uphills also depend on the powertrain.

Hills impact vehicle energy efficiency through the following mechanisms:

- **Braking Inducement:** Hills can force a vehicle against some constraint, such as a speed limit, or for platoons, a desired gap. When hills result in the application of negative powertrain torque (active deceleration), energy is wasted [10].
- **Energy Source/Sink:** Hills lead to a change in total energy, either taking vehicle energy to go uphill or returning energy on downhills, but this in itself can either dilute or strengthen platooning benefits[5].
- **Powertrain Map Shifts:** Hills also cause varied power demands, changing the powertrain operating regime. Hill-induced downshifts, for example, usually change vehicle energy efficiency.

As a result of these effects, fuel-efficient driving over hills has the primary objective of minimizing active deceleration, with the secondary aim of minimizing suboptimal powertrain operation. This applies to all vehicles, including platoons.

For both single vehicles and platoons, eco-driving over hilly terrain is typically accomplished using look-ahead control [60]. There are some examples of eco-driving cruise controllers that do not use a look-ahead [61]. Many commercial cruise controllers simply allow for deviation in the set speed (referred to as flexible cruise control), but hilly terrain is enough of a challenge for HDV fuel consumption that all major truck manufacturers offer some form of look-ahead cruise control[6].

[5] As a consequence, it is important to avoid comparing results with different net elevation changes [59].

[6] All major heavy duty truck manufacturers offer predictive cruise control (PCC) that uses upcoming grade to adjust velocity accordingly [62]. A step under PCC is flexible cruise control (FCC). FCC allows for over/undershoot of a set speed on hills, also called droop. In one experiment, the FCC reduced fuel consumption compared to standard CC, although even greater savings were achieved by a prospective PCC [63].

Real roads are rarely composed of constant uphill or downhill sections. Often, crests and valleys occur as in Figure 2.8, putting the leader and follower under opposite-sign grade forces.

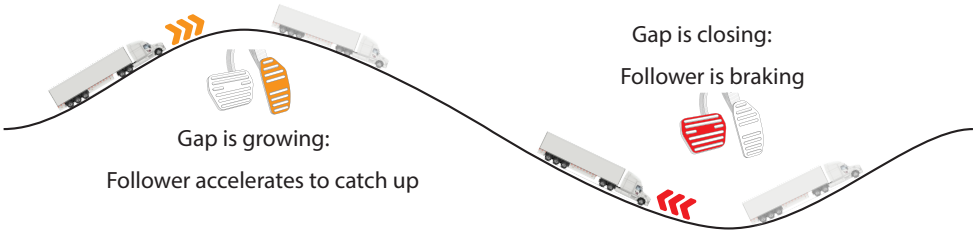

Figure 2.8: Platoons traversing grade are faced with disturbances that are at odds with strict longitudinal spacing.

When speed limits are in effect, the fuel-optimal strategy for traversing hill crests is to preemptively decelerate before the top of the crest is reached, thus reducing the amount of braking required on the steep downhill tract [50].

The impact of simultaneous shifting and look-ahead control of trucks has been explored in [64, 65, 63]. Simultaneous shifting was determined to produce tighter gap control and driver comfort, but look-ahead control was ultimately required to maintain platooning aerodynamic benefits over challenging grade.

Regarding mass variations, placing a heavier vehicle in the back increases the likelihood of hill steepness and may induce situations where the platooning benefit is compromised on downhills [60], as well as situations where the heavy follower gets "left behind" on uphill sections. In contrast, a heavy leader could constrain the follower to its limits for uphill steepness.

Fuel-efficient platooning over grade is a challenging research topic, and there are still many open research questions regarding efficient platooning over hills. What is known is that hills can seriously compromise platoon energy benefits and that effective controller design is key. Attention now shifts from hills to platoon interruptions caused by other vehicles and roadway obstructions.

2.5.2 The Effect of Platoon Interruptions

Platoon interruptions are defined as traffic disturbances that force an operational change. These interruptions come in two main forms:

- Cut-ins by other vehicles
- Speed changes, due either to congestion, or roadway construction

The aerodynamic benefits of platooning under traffic disturbances can be treated as quasi-steady [66]. Platoon interruptions affect energy usage through active deceleration, just as hills do [67].

There are only a few focused efforts on the effect of cut-ins on a platoon's fuel economy [26, 42, 68, 69]. Cut-ins in [26] slightly impacted the fuel economy of the platoon using California PATH's control strategy, but for the Auburn University PID-based CACC system in [42], cut-ins completely erased platoon benefits.

Other cut-in analyses aim at naturalistic characterization of cut-ins [70, 71, 72, 73, 74], prediction, and modeling of cut-ins [75, 76, 77]. Cut-ins are unlikely to occur at very close follow-up distances. There is a separate class of work that discusses how a cut-in shall be handled, which is somewhat relevant here. See [78] for an example that discusses the logic for handling cut-ins.

The impact of cut-ins on a platoon's fuel consumption is still nebulous. It remains an open research question, and there is no consensus on how to allocate the energy losses resulting from a cut-in.

Regarding the impact of upcoming traffic disturbances on a platoon, the literature is focused on the wider traffic benefits of platooning. High-level speed planners can handle upcoming speed changes, provided that this information is available [79, 80]. There are some indications that smart vehicle technologies, including platooning, carry system-wide benefits, but this remains to be validated in practice [81].

In summary, there is no clear research understanding of the impact of cut-ins and traffic disturbances on platoon energy consumption. In particular, cut-ins have the potential to reduce leader benefits and erase follower benefits, but this is not always so. The next section will cover the effect of terrestrial winds and other traffic wakes on platoon performance.

2.5.3 The Effect of the Terrestrial Environment and Other Naturalistic Traffic Wakes

The presence of terrestrial winds and other traffic induces changes in the aerodynamics of road vehicles [82, 83]. A direct headwind adds additional aerodynamic drag, and a direct tailwind reduces it. The most difficult case to model is when the wind is at an angle to the direction of a vehicle, causing a yaw. However, the yaw-induced wind angle from terrestrial winds alone is typically less than 9 degrees [43][7].

First, we examine how wind-induced yaw affects platooning. Drag reduction is not erased by yaw conditions [25], but aerodynamic devices (skirts, boat tails) typically lose some of their effectiveness with yaw [86]. As usual, lead vehicles in a platoon are less sensitive to yaw conditions than the following vehicles: An aligned platoon shows less than 5% difference for the lead vehicle energy savings at a 15° yaw angle [32]. This 5% difference is not from a loss of the base pressure effect, which is relatively insensitive to yaw. Instead the reduction comes from the dilution of benefits: as yaw increases the overall vehicle drag, the "push" effect is diluted.

As for the trailing vehicle, there are three ways that yaw impacts the aerodynamic drag according to [16]:

1. As the wake shifts due to yaw, stagnation pressure and the flow acceleration on the exposed corner increase, which increases drag,
2. A lateral offset into the shifted wake decreases the effective yaw angle, which decreases drag.

[7]Several techniques to adjust a road vehicle's zero-yaw drag for roadway conditions were compared in [84], including the SAE 1252 procedure [85].

3. The overall effective yaw angle is reduced due to the decreased overall wake, which decreases drag.

As [16] points out, the overall effect of yaw on the trailing vehicle is an increased drag at small separations (5 m) and a decreased drag at larger separations (20 m). Behaviorally, the effective yaw experienced by the trailing vehicle is less than the free-stream yaw and may even reverse at inter-vehicle distances of less than 30 meters (100 feet) [33]. Due to this reduced effective yaw angle, the drag reduction may be more than 40% higher at 4° than at 0°.

To summarize the effect of wind-induced yaw on platoon drag, the effect of wind-induced yaw on a platoon is only dilutional for leaders. For followers, yaw has much more complex effects. In particular, at angles of yaw lower than 5°, the followers are less sensitive to yaw than isolated vehicles, reducing the penalty of yaw on energy consumption.

Now that the effect of terrestrial winds on platoons has been reviewed, we will shift focus to the effect of traffic wakes and terrain obstructions on the aerodynamics (and therefore savings) of platooning. The measurement and classification of roadway turbulence was carried out in [87, 88, 89]. These on-road measurements can be used to represent realistic turbulence in simulations and wind tunnels [90, 91].

The presence of traffic provides some aerodynamic benefits to otherwise isolated vehicles [92]. This benefit from surrounding traffic has been termed naturalistic or background platooning. Vehicles ahead in the same lane provide a slight amount of aerodynamic reduction (relative to platooning reduction). Vehicles ahead in an adjacent lane also reduce aerodynamic drag with a shifted wake. While trucks are rarely in a natural platoon with other trucks (97.8% of the time trucks are > 90 meters from one another, [18]), the question of how much naturalistic platooning is already occurring is of vital importance. Test track results have suggested that the benefits of platooning are mostly additive to the benefits of background platooning [42].

In summary, terrestrial winds and other traffic affect the benefits of trailing trucks more than they affect the benefits of lead trucks. Wind-induced yaw angles are unlikely to exceed 9° on the road and will often be much less than that.

2.6 Summary of Various Effects on Platooning Energy Savings

To this point, platooning energy savings have been reviewed comprehensively. Before moving on to existing models of platooning energy savings, a summary is warranted.

- As the longitudinal spacing increases (larger IVD), platoon savings decrease nonlinearly for both leaders and followers.
- As lateral offset increases, the platoon benefits decrease (nonlinearly).
- As speed increases, potential benefit increases (concentration).
- As mass increases, potential benefit decreases (dilution).
- Placing a more aerodynamic truck/trailer in the front of the platoon is more advantageous than placement elsewhere.
- Pure steep downhills nullify platooning benefits, and uphills dilute benefits proportionally to grade.
- Cut-ins dilute leader benefits, but cause unbounded energy losses for followers.
- Yaw angle dilutes leader benefits in proportion to the isolation drag increase but increases benefits for followers due to decreased sensitivity.

Figure 2.9 further summarizes the marginal effects of internal and external factors on platoon fuel consumption. Some of the factors that could not be included in Figure 2.9 are

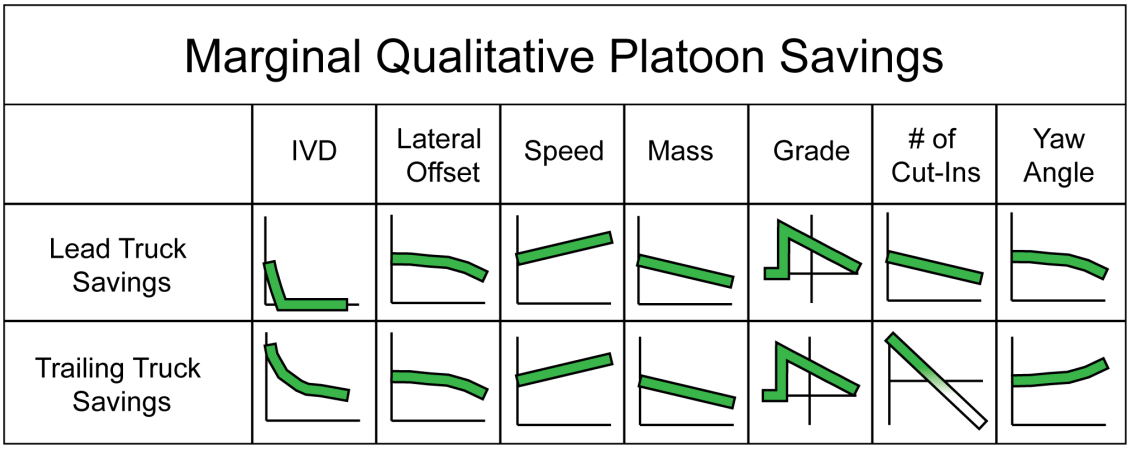

Figure 2.9: Marginal effects of various factors on platoon energy savings. Savings are relative to vehicles in isolation with the same factor levels.

the influence of truck and trailer styles and the influence of time-varying grade. Furthermore, unidirectional powertrain flow is supposed for the grade and cut-in trends, leading to zero/negative savings, but regenerative braking would have different trends. The combined effect of all the factors mentioned above is complex. The factors of grade and cut-ins, for instance, are a strong (and presently unknown) function of the gap control strategy, the longitudinal spacing, the speed, and the mass of the vehicles. Representation of such complex interactions is a modeling task, so next we turn our attention to the state-of-the-art platoon energy prediction models.

2.7 Existing Models of Platooning Energy Savings

Platooning energy savings models seek to unify the effects of internal factors and disturbances to either predict drag reduction or actual fuel savings. Applications of platooning energy savings models include:

- Simulation of platooning benefits in the larger traffic simulation
- As input to a cost function for platooning optimization and control

Several types of platoon energy prediction models exist. A fully empirical model of platoon drag savings for LDVs, HDVs, and buses may be found in [93]. The model is a rational polynomial, which gives the drag reduction ratio $\frac{C_D}{C_{D\infty}}$:

$$\frac{C_D}{C_{D\infty}} = \begin{cases} \frac{a_N G^N + a_{N-1} G^{N-1} + ... + a_1 G^1 + a_0}{b_N G^N + b_{N-1} G^{N-1} + ... + b_1 G^1 + b_0}, & 0 < G < G_o \\ 0, & G \geq G_o \end{cases} \tag{2.1}$$

G is the separation gap, G_o is the gap at which no benefit occurs, and $a_{N:0}, b_{N:0}$ are empirically determined coefficients. A method to account for crosswinds is incorporated, which makes use of both drag and lift coefficients. This model is converted to a fuel curve using a road

load model [94]:

$$P(t) = \left(\frac{R(t) + 1.04ma(t)}{3600\eta_d} \right) v(t) \tag{2.2}$$

$$R(t) = \frac{\rho}{25.92} C_d C_h A_f v(t)^2 + Gm \frac{C_r}{1000} (C_1 v(t) + C_2) + gmG(t) \tag{2.3}$$

A piecewise second-order polynomial fit is applied to convert power at the wheels to fuel usage, with a constant loss at negative power.

$$F(t) = \begin{cases} \alpha_0 + \alpha_1 P(t) + \alpha_2 P(t)^2, & P(t) \geq 0 \\ \alpha_0, & P(t) < 0 \end{cases} \tag{2.4}$$

The form of the model in [94] would capture controller-induced losses were there a mechanism in it to separate the controller-induced braking from other reasons for braking.

Beyond purely empirical models, there are three hybrid prediction models that blend physics with empiricism [27, 28, 95]. [27] is the most general formula, as well as the basis for [28, 95]. Base pressure drag reduction and wake deficits were measured in a wind tunnel experiment with passenger car models, and wake theory was leveraged to generalize the results for platoons of differently sized vehicles. The drag reduction ratios $\frac{C_D}{C_{D_\infty}}$ for the n^{th} vehicle are given as a function of the drag reduction ratio of the preceding vehicle. The model is only meant for aerodynamic drag reduction prediction, not holistic power/energy use.

[28] is a development of [27] for HDV platoons. It, too, is only concerned with the drag reduction of platooning. It was found that there is a mistake in the published coefficients of [28], but there was enough information to reconstruct them [96]. Those reconstructed coefficients are presented in Table 2.2.

Finally, [95] further develops [27, 28] to include the effects of IVD, lateral offset, and crosswinds.

Table 2.2: Original published values for the Schmid et al. drag reduction model [28] compared to the recalculated values [96].

	X_1	X_2	X_3	X_4	X_5	X_6
Published	0.1821	1.5216	0.6643	8.9340	0.3371	-0.0419
Recalculated	5.4543	1.5197	0.6610	8.9289	0.3374	-0.0422

A different tier of models uses an offline CFD model of drag reduction coupled with machine learning methods to enable real-time estimation of platoon drag coefficients. Notable examples include [97, 98, 99].

In the context of braking loss models, various researchers have highlighted the significance of braking losses in platooning, but few have quantified this loss experimentally; none have introduced a specific "cost" associated with braking [60, 15, 67]. The closest the literature has come to this is in [67], where friction braking work was estimated for the experimental data using an OEM-developed black box. Some work has assumed that the brake pressure is linear with the brake force [100, 101]. There is a lack of available techniques in published literature for determining the precise braking force encountered by a vehicle. Hence, it can be concluded that existing research has not presented a method for extrapolating the advantages of platooning in braking scenarios, apart from the author's own research.

In the work of the author, practical methods are given to calculate the braking power with vehicle acceleration and boolean braking status only, using empirical fits. The primary emphasis of these studies lies in quantifying braking rather than drag, although both aspects are acknowledged. In [10], a method is developed to obtain mechanical energy losses due to controller-induced braking without any knowledge of the magnitude of the braking force. The active deceleration energy (E_{AD}) of the heterogeneous platoons of two and four trucks with various controllers is calculated, giving an estimate of the increase in effective fuel use versus E_{AD}. In [11], the previous work is further developed with a simulation analysis and extension of the E_{AD} calculation to real-time.

2.8 Research Gap and Motivation

And so it seems that the reduction of controller-induced losses is of great importance to platooning efficiency, but it has not yet been included. Except for the work of the author, all platoon energy prediction models are built for aerodynamic drag reduction or comprehensive energy use. So far, no model framework has been provided to incorporate the effects of controller-induced losses along with the aerodynamic benefits of platooning, which is a first-principles-based approach. Furthermore, no research has demonstrated a prediction model that can operate in situ and provide feedback to the operator of a platoon vehicle. With this research gap in mind, we will proceed to describing the proposed methodology for real-time platoon energy savings estimation.

Chapter 3

Methodology

In the previous chapter, it was revealed that there is a significant gap in understanding the differential energy consumption of platoons. Although there are several platoon drag reduction models and separately there are comprehensive fuel consumption models, there are no models that presently combine the aerodynamic benefits of platooning with the unbounded energy losses due to actuation effort. In this section, such a model is proposed.

First, we treat in depth the physical principles that govern the problem. Then, the author's previous work on braking loss characterization is seen to not generalize to other data. This motivates a new approach. High-fidelity platooning simulations are used to formulate the framework in a way that is valid across a wide range of speeds, vehicle masses, and drag reduction ratios. Following the development of the framework is a section devoted to the practical implementation of the framework that uses nothing more than readily available vehicle CAN signals. Finally, the chapter concludes by describing the validation approach, including a description of the data that will be used for the validation.

3.1 Platoon Energy Consumption Fundamentals

To properly formulate a model for differential platooning energy consumption, we must start with the first principles. The energy flow pathways for all ground vehicles can be understood in Figure 3.1. As road conditions and tractive requirements change, the amount of energy flowing on the pathways in Figure 3.1 changes continuously. Vehicle energy flow begins at one or more sources, be it liquid fuel, batteries, or both. The energy is converted into useful work by the powertrain. During this process, some energy is also lost as waste heat, but the split between useful work and waste heat varies widely with the powertrain architecture and operating conditions[1]. The useful work may then be applied in three ways:

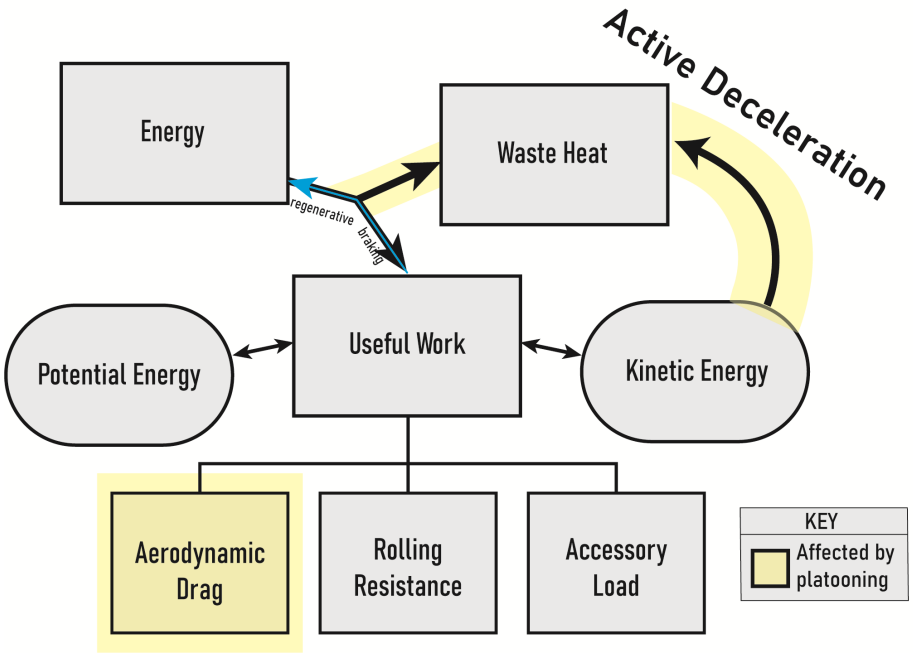

Figure 3.1: Energy use flowchart for vehicles.

1. to increase the vehicle's potential energy, such as is needed to climb a grade.
2. to increase the vehicle's kinetic energy, resulting in a higher vehicle velocity.

[1]For internal combustion powertrains, conversion efficiencies are often about 20% and rarely exceed 40% in the best conditions. For battery electric powertrains, approximately 90% of the battery energy is converted to useful work.

3. to overcome **road loads**, namely aerodynamic drag, rolling resistance, and accessory load.

Once the kinetic energy or potential energy is positive, it can also be reversed and used to overcome the road loads. Active deceleration can be utilized to actively slow down the vehicle. For conventional brakes, this means that kinetic energy is directly converted to waste heat. For vehicles with regenerative braking, some of this is reclaimed, perhaps 50% [102].

The yellow-shaded regions in Figure 3.1 depict the energy flow pathways that are affected by platooning: (i) aerodynamic drag, (ii) active deceleration, and (iii) waste heat losses[2]. Energy lost to aerodynamic drag, which is the target of the platoon strategy, is reduced by platoon operation. Regarding active deceleration, maintaining the platoon gap generates additional braking in the presence of disturbances, which leads to a flow of kinetic energy to waste heat. The negative impacts of this behavior can be partially mitigated by using a regenerative braking strategy. Finally, platooning impacts waste heat output in two ways:

1. The overall road load is reduced through aerodynamic drag changes, which may change the conversion efficiency of the powertrain.
2. The controller-induced dynamics have the potential to move the operating point on the powertrain efficiency map. For example, to catch up after a braking event, a downshift may be required, leading to lower efficiency.

In summary, all vehicles share common energy flow characteristics. Platooning affects three portions of the energy flow directly: Aerodynamic drag is the source of platooning benefits, active deceleration is the source of platooning losses, and waste heat is a complicating factor. To characterize the energy usage of the platoon, it is necessary to consider the importance of drag reduction and active deceleration in relation to other factors that remain unaffected by the platooning process.

[2]It is possible that the accessory load losses may be impacted by platooning. At small IVDs, where cooling airflow is reduced, an engine fan may be required.

3.1.1 Model Assumptions

The proposed framework seeks to infer the platooning energy benefits during a platooning microtrip, which is defined as the duty cycle from the time a team of trucks engages in platooning to when they disengage. The platooning microtrip may contain cut-ins with other vehicles, hills, and mild speed changes, but not complete stops. More specifically, six assumptions are made about the platooning energy consumption over a given microtrip:

1. The first assumption is that **differential platoon energy consumption is driven by the balance of drag reduction and braking energy consumption, whereas other road loads simply dilute the benefits of platooning**[3]. There is one caveat: when platooning does not alter the velocity trajectory of a platoon member, then only the aerodynamic benefits count towards the platoon's energy consumption. A lead vehicle that uses ACC falls under this exception.
2. The second assumption is that **the efficiency of the powertrain is not changed significantly by platooning**. Another way to state this assumption is that the power at the wheels is proportional to the energy consumption of the powertrain. A byproduct of this assumption is that fuel use is proportional to mechanical power output. This linear assumption is often made in a model-predictive control context [48, 67]. This assumption is made for two reasons:

   (a) it preserves generalizability across various powertrains

   (b) a validated, high-fidelity powertrain efficiency map was not available

   The assumption of constant powertrain efficiency is a simplifying one which will often be violated in reality, depending on the situation. To avoid making this assumption, one could include a powertrain efficiency map that maps the power output at the wheels

[3] To illustrate, although both increased braking effort and increased vehicle mass would lead to a decrease in percent savings, only the increase in braking can be attributed to platooning. To put it another way, there is a situation in which the average braking power could increase enough that platooning consumed more fuel than non-platooning, but increasing mass would only ever reduce the potential benefits ("dilute"), never directly counteract them.

to the power / energy input to the powertrain. By doing so, the nonlinear, time-dependent behavior would be accounted for, increasing the precision of the platooning benefit estimate. This is left to future work.

3. The third assumption is that **the energy lost to braking is manifested in the power consumption of the platooning truck during the microtrip**. This assumption ignores this time dependence by modeling the braking loss as an instantaneous increase in energy consumption. While a truck is actively decelerating, it consumes less fuel than if it were to stay on throttle. Only after some time has passed does the braking energy loss lead to powertrain energy consumption.
4. The fourth assumption is that **the non-platooning road load parameters are known, such that the power consumption of the truck can be modeled.** At the very least, knowledge of these road loads will require the truck mass (m), the aerodynamic drag coefficient (C_d), and the rolling resistance coefficient (C_{rr}). Under this assumption, the model is subject to incorrect parameterization. To assess the validity of this assumption, a sensitivity analysis to road load parameterization will be conducted in Section 4.1.
5. The fifth assumption is that **the mean values for drag reduction, speed, and braking power are representative of the mean differential energy consumption of platooning**. Under this assumption, the running microtrip averages can be used to estimate energy benefits.
6. The sixth assumption is specific to cut-ins: **the braking losses from a vehicle cutting in between the platoon are a platoon-induced loss, not a diluting effect.** Treating a cut-in as a dilution loss would assume that the cut-in vehicle maneuvers the same way when cutting in front of a single vehicle as it does when cutting between two vehicles. Treating cut-ins dilutionally seems too relaxed: after all, a vehicle that cuts in between a platoon has no space to continue forward and must match the speeds the leader. The induced braking losses would likely be reduced if there were space for

the cut-in vehicle to continue forward. By assuming that braking losses are due to platooning, rather than a contributing factor that dilutes benefit, the model is more conservative about platooning benefits.

Under these six assumptions, a platoon's energy consumption relative to a non-platooning reference is a function of energy lost to braking and energy saved by drag reduction. The powertrain efficiency is assumed constant whether platooning or not, which makes fuel and mechanical power proportional. It is assumed that the road loads are well-parameterized and that the submodels for drag reduction and energy lost to braking take the correct form.

An assumption that was believed to be essential at first, but turned out not to be, is that the potential energy is equal at the start and end of the platooning microtrip. In other words, there is no net elevation change between the start and end. This assumption of zero net potential energy change is not necessary, but does have some interesting effects on differential platooning energy consumption which are covered in Appendix B.

The key assumptions of the framework have been formalized in this section. There is still the question of what the platoon energy consumption is relative to, a question which will be answered next.

3.1.2 Choosing the Non-Platooning Reference

The general framework to infer the change in platooning energy consumption is shown in Figure 3.2. These inferences are represented in Figure 3.2 in terms of both power ($NPC_{inferred}$ and $\Delta P_{inferred}$) and fuel ($NFC_{inferred}$ and $\Delta F_{inferred}$). The platoon energy in-

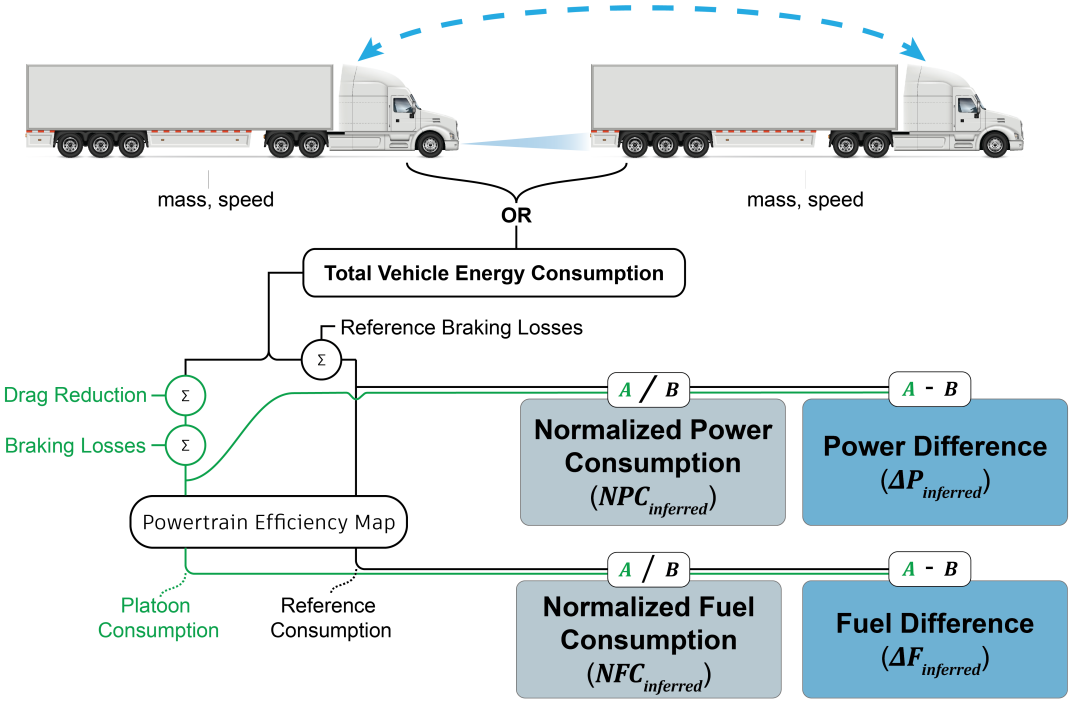

Figure 3.2: The top-level architecture of the proposed platooning energy consumption model.

ference involves some theoretical **non-platooning reference energy consumption**, which requires further definition. Referring to Figure 3.2, what reference energy consumption belongs to the term marked "B"?

Some desired characteristics of the reference energy consumption, which would occur over the corresponding platooning microtrip, are:

- The same average speed or trip time from start to end
- The same initial and final speed[4]
- The same environmental factors and operating conditions

[4]This condition becomes less consequential as the trip becomes larger and the total energy consumed by road loads overwhelms the kinetic and potential energy requirements.

These characteristics of the reference energy consumption are not to be confused with the assumptions made about the platooning energy consumption in the preceding subsection.

If these characteristics about the reference consumption are reasonably satisfied by two or more given duty cycles, then confident conclusions can be drawn about the relative efficiency of duty cycles[5]. In constant-speed test track experiments, it is usually assumed that these assumptions are satisfied, which is how the platooning energy benefits have been previously ascertained. In contrast, there will not necessarily be a reference real-world platooning microtrips. Therefore, it is difficult to assess the real-world platoon energy consumption on a relative basis.

The ideal reference for platooning energy consumption would be an identical vehicle with all things equal except for those factors influenced by platooning. This is impossible experimentally, as it would require a non-platooning twin of the truck shadowing it. To get around this, in prior platooning experiments, it has been assumed that either:

- The experimental conditions are relatively constant, and baselines have been collected at a different time.
- The ratio of performance between the truck and a control truck follows the same trend as the truck would have if it could "shadow" itself in a non-platooning configuration.

The SAE J1321 Type II fuel test assumes the latter when it defines the percent benefit using the test-to-control fuel consumption ratio (the T/C ratio).

The question of reference is further complicated by the inclusion of braking losses. Braking is an unbounded loss that is determined by the driver or the longitudinal dynamics controller. Two choices of reference energy consumption are proposed for incorporating braking in the context of platoon benefits:

[5]There is one caveat: two trips with the same travel time and boundary conditions do not necessarily have the same mean aerodynamic load. Assuming that drag is a function of velocity squared, it can be shown that the mean aerodynamic load increases proportionally to the squared mean velocity plus the velocity variance: $\bar{P}_{aero} \propto \bar{v}^2 + \sigma_v^2$. At highway speeds, the velocity variance σ_v^2 is four orders of magnitude less than the squared mean velocity $\bar{v}^2$. Therefore, we may assume that velocity variance increases aerodynamic load negligibly.

1. **Brakeless**: Compare platooning energy consumption to a brakeless reference that accomplishes the same mission (i.e. trip time, mean speed, boundary conditions).

OR

2. **Realistic**: Compare the platoon to a realistic longitudinal velocity reference that allows braking (e.g. ACC) and accomplishes the same mission.

There are advantages and disadvantages to both a brakeless or a realistic reference.

Referencing a theoretical brakeless scenario is attractive because it is void of all braking losses, providing a reliable zero. Furthermore, it is simpler than a realistic reference because it is only a function of mass and speed, not the amount of braking that occurred. However, on its own, the use of a brakeless reference may lead to the incorrect conclusion that platooning does not save fuel when braking is unavoidable even without platooning. Finally, there is the possibility that a brakeless reference may be unrealistic or even impossible. An example of this is a platooning microtrip on a long steep downhill, where all vehicles must brake for safety and legal reasons.

On the other hand, a realistic reference allows for the possibility that some amount of active deceleration is unavoidable. Thus, a realistic reference is more optimistic about platooning energy consumption than a brakeless reference. Although truck cruise control systems are designed to handle small overshoots in speed without braking, there are situations where speed limit constraints sometimes require braking. It would be incorrect to call braking losses in such a situation platoon-induced losses. In fact, using a realistic velocity reference, a platooning truck that braked less than the reference would benefit from more than just aerodynamics. In other words, platooning can sometimes help avoid braking losses, as in the ACC characterization in [60]. The realistic reference could be defined by using the leader's velocity behavior for all of the follower's reference. This "leader-referenced" approach may work if the trucks are identical and the leader is using a disconnected form of cruise control, but if the masses and base aerodynamics are different, it is unclear whether

this approach will work. Additionally, the amount of braking in the realistic reference varies unpredictably and is specific to how the realistic reference is generated.

To summarize, it is best to have both a brakeless reference and a realistic one to compare a given platooning result with and use the brakeless reference as a control and the realistic reference to make practical conclusions about whether platooning saved fuel or not. In either case, a suitable reference must have the same travel time and boundary conditions as the platooning cycle in question. Next, the calculation methodology for the platoon energy consumption will be described.

3.1.3 Assessing the True Platooning Energy Change

While the assumptions and philosophy of the model have been covered, the calculation of the true energy consumption difference of the platoon has not yet been introduced. The purpose of this section is to show how the actual difference in platooning consumption can be calculated.

The average reference power/fuel consumption as a function of the vehicle mass m and the mean speed $\bar{v}$ is defined by Equations 3.1 and 3.2:

$$\bar{P}_{ref}(m, \bar{v}) = \frac{1}{\Delta t} \int P \, dt \mid (P_{AD} = P_{AD,\, ref},\ DRR = 1) \tag{3.1}$$

$$\bar{F}_{ref}(m, \bar{v}) = \frac{1}{\Delta t} \int F \, dt \mid (P_{AD} = P_{AD,\, ref},\ DRR = 1) \tag{3.2}$$

which is the average of power consumption (P) or fuel consumption (F) given a reference level of braking ($P_{AD,\, ref}$) and zero drag reduction ($DRR = 1.0$). As discussed previously, ($P_{AD,\, ref}$) will be zero if using a brakeless reference.

To express the true difference between platoon power consumption and reference power consumption, we can work in either absolute or relative terms:

$$\Delta P_{true}(m, \bar{v}) = \bar{P}_{plat}(m, \bar{v}) - \bar{P}_{ref}(m, \bar{v}) \tag{3.3}$$

$$\Delta F_{true}(m, \bar{v}) = \bar{F}_{plat}(m, \bar{v}) - \bar{F}_{ref}(m, \bar{v}) \tag{3.4}$$

$$NPC_{true}(m, \bar{v}) = \frac{\bar{P}_{plat}(m, \bar{v})}{\bar{P}_{ref}(m, \bar{v})} \tag{3.5}$$

$$NFC_{true}(m, \bar{v}) = \frac{\bar{F}_{plat}(m, \bar{v})}{\bar{F}_{ref}(m, \bar{v})} \tag{3.6}$$

where ΔP_{true} and ΔF_{true} are the absolute power and fuel difference, and NPC_{true} and NFC_{true} are Normalized Power and Fuel Consumption. The subscript "true" distinguishes them from "inferred", which will be developed later. We can use absolute and relative expressions equivalently to make statements about platoon fuel consumption, although we will prefer to work with NPC_{true} and NFC_{true}, since it allows a direct comparison of differential energy consumption values from different trucks.

Because NPC_{true} and NFC_{true} are ratios relative to a reference consumption, they can be interpreted as a percent benefit. For example, $NPC_{true} = 0.9$ could be interpreted as "the platooning run consumed 0.9 times as much power as the reference run (90%)".

Under the constant powertrain efficiency assumption (Model Assumption 2), NPC_{true} and NFC_{true} are equivalent:

$$NPC_{true} \approx NFC_{true} \tag{3.7}$$

Figure 3.3 gives evidence for Equation 3.7 as NPC_{true} and NFC_{true} mostly fall along a 1:1 line. The nonlinearity in the simulated powertrain efficiency manifests itself as an upward shift from the 1:1 line, particularly in the region below $NFC = NPC = 1$. Figure 3.3 is the result of a high-fidelity vehicle simulation, which will be described in Section 3.3.1.

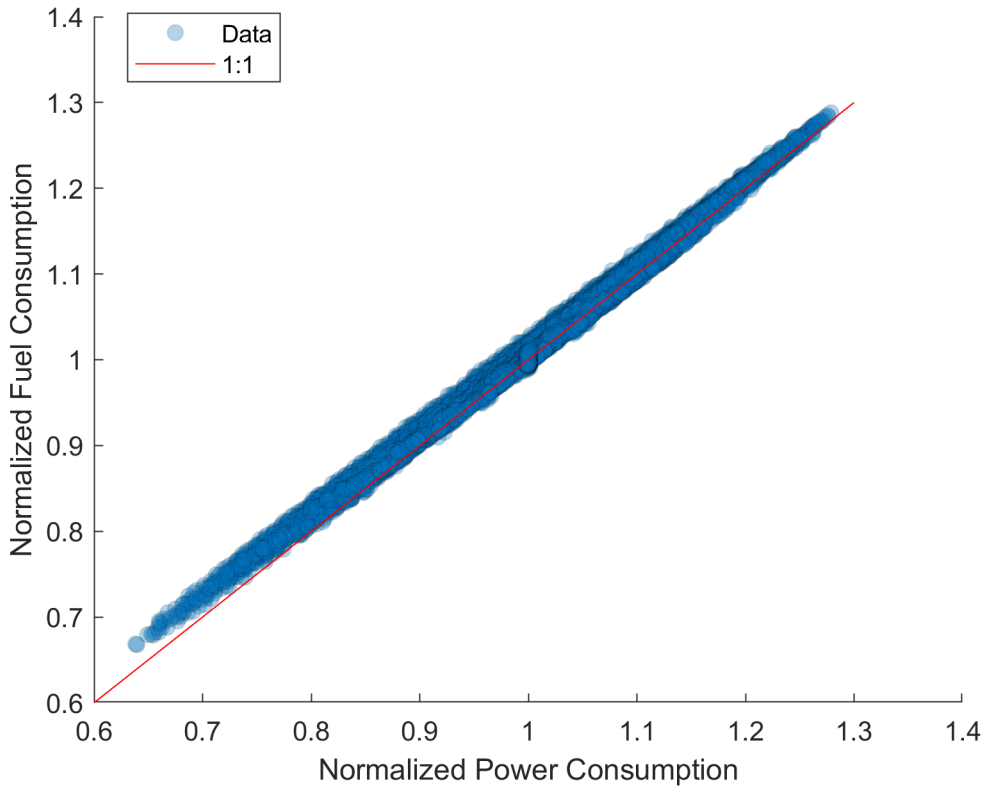

Figure 3.3: Normalized Power Consumption versus Normalized Fuel Consumption. Results from simulation as described in Section 3.3.1 and Appendix C.

In the upcoming section, we will look at how the already published approaches have inferred NFC_{true}. These methods are promising, but their applicability may be limited.

3.2 Assessing the Adequacy of Available Methods

It is supposed that as NPC_{true}[6] can be inferred as some function of drag reduction ratio, mass, speed, and braking losses:

$$NPC_{inferred} = f(DRR, P_{AD}, m, v) \tag{3.8}$$

The form of Equation 3.8 has been speculated on in two previous publications by the author [10, 11]. The energy lost to braking was formally termed Active Deceleration Energy (E_{AD}). It was shown that E_{AD} was linearly related to the change in NFC_{true}. The linear relationship was described empirically by Equation 3.9:

$$NFC = \beta_0 + \beta_1 \frac{E_{AD}}{\Delta t \cdot m_{eff}} + \varepsilon \tag{3.9}$$

where Δt it the trip time of the platoon, and m_{eff} is effective linearized mass of the truck to include the rotating mass contribution, and the β's are least-squares fit coefficients.

The methods in those works linking E_{AD} with NFC_{true} did well to relate braking to the energy consumption of truck platoons. However, in those publications, the trucks were tested only at 45 mph (20.1 m/s), and the trucks operated only unloaded. As such, previous efforts to correlate drag reduction and braking with energy change may not apply to platoons of varying mass and speed. Before going further, we will have a brief review of the previous methods.

3.2.1 Brief Review of Previous Methods to Infer Benefits

The first paper [10] introduced E_{AD} and derived a batch post-processing procedure to estimate the total braking losses from the experimental data. The routine is depicted in Figure 3.4, which shows the segregation of braking and coasting in Figure 3.4a, and the

[6] and equivalently, NFC_{true}, though there will be a term to convert mechanical power to fuel power

resulting relationship between E_{AD} and NFC_{true} in Figure 3.4b. To account for not only

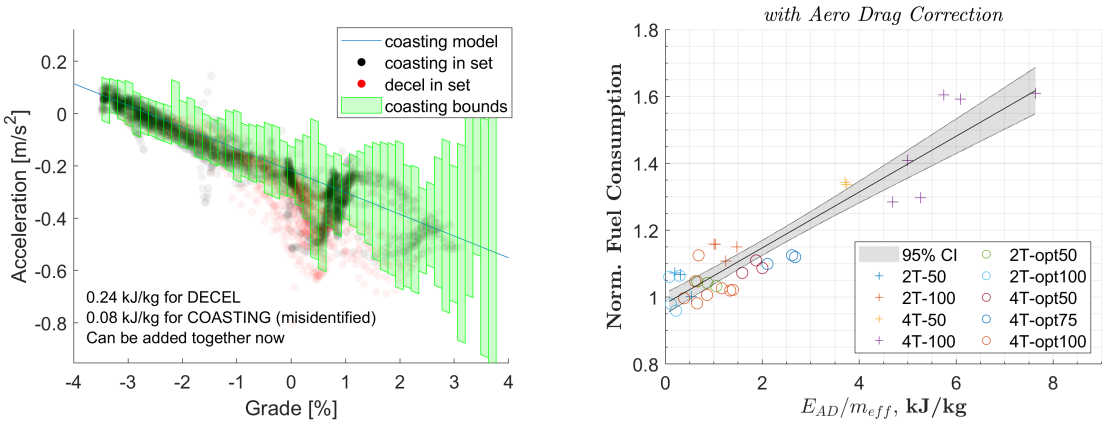

(a) A demonstration of how braking acceleration was separated from coasting acceleration.

(b) The resulting relationship between E_{AD} and NFC, adjusting for drag reduction using the "drag is half the road load" assumption.

Figure 3.4: Figures from [10] demonstrating the previously published method for relating braking to NFC.

the braking but also the drag reduction of the platoons, the power law drag reduction ratio of [93] was applied to NFC_{true}. This drag reduction ratio was applied in two different ways: (1) by assuming that drag represents half of the vehicle's road load, or (2) by treating the drag reduction ratio as a regressor for NFC. In general, the results in [10] indicated that E_{AD} described the losses induced by the platoon controller, but that it was difficult to distinguish coasting deceleration from active deceleration, which affected the estimate of E_{AD}.

The second paper introduced a way to estimate E_{AD} in real time, in an effort to increase the practical applicability of the method. Where the first paper was a proof-of-concept, the second paper was more concerned with a practical implementation of the theory that NFC is linearly related to braking and drag reduction.

Using high-fidelity simulation in the virtual test environment software IPG TruckMaker[7], the method in the second paper was found to give an E_{AD} within 5% of the true value, assuming a true knowledge of vehicle parameters. The panels in Figure 3.5 summarize the message of the second paper, [11].

[7]Please refer to the original source [11] for the details of the simulation in that work, which is nearly identical to the simulation environment described in Section 3.3.1 and Appendix C.

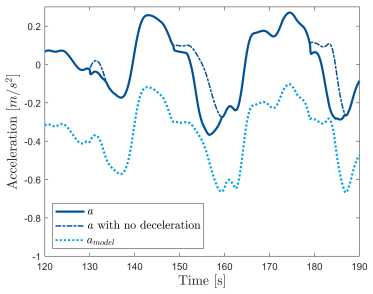

(a) True acceleration (solid dark), modeled brakeless acceleration (dashed light), and theoretical brakeless acceleration (dashed dark).

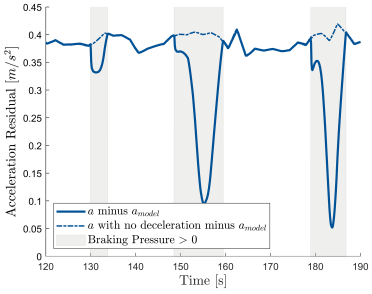

(b) The residual of true acceleration and modeled acceleration. Gray regions represent when the vehicle was braking.

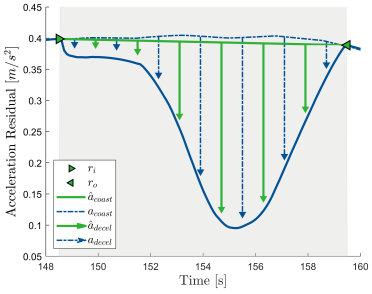

(c) Estimated acceleration residual (green) versus the true acceleration residual (dark blue).

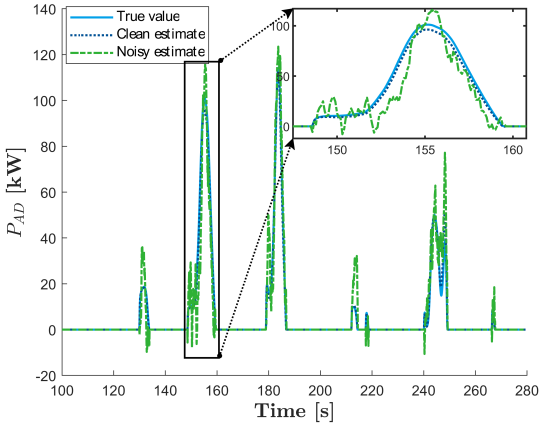

(d) A demonstration of the estimation of P_{AD} in simulation.

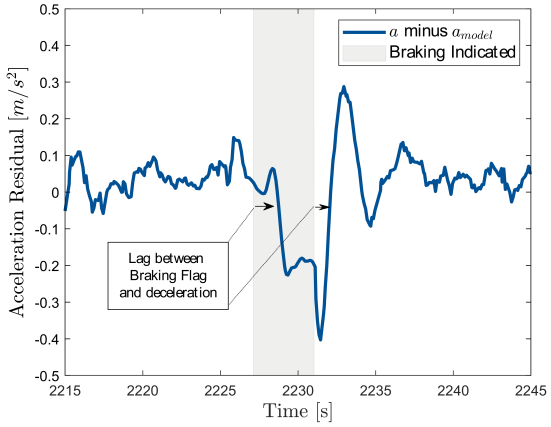

(e) Delays in the indicated braking status and the actual deceleration.

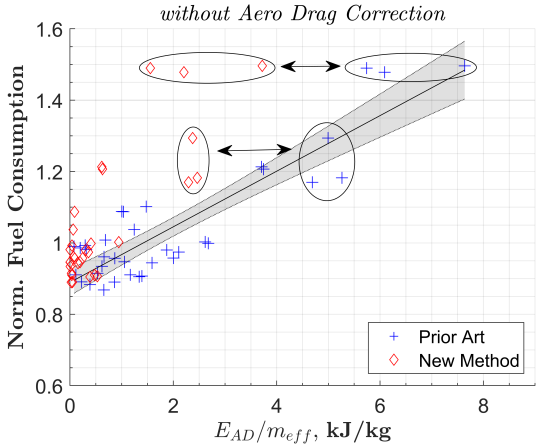

(f) Results from papers [10] ("Prior Art") versus [11] ("New Method") without accounting for the delay in Figure 3.5e.

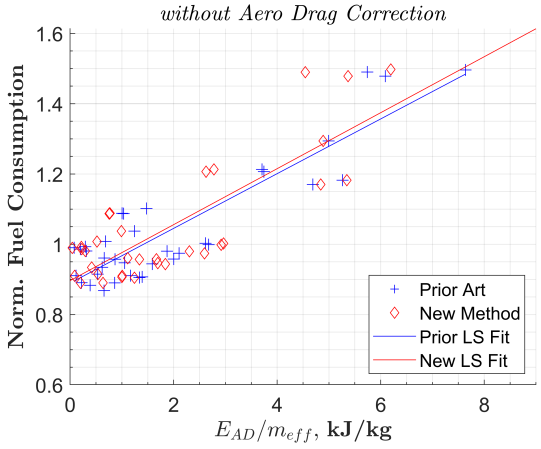

(g) Results from papers [10] ("Prior Art") versus [11] ("New Method") after a rough attempt to account for the delay in Figure 3.5e.

Figure 3.5: Figures from [11] demonstrating the previously published method for relating braking to NFC.

The first step in the calculation process for E_{AD} is to acquire an acceleration measurement and separately model the acceleration of the vehicle without braking, as shown in Figure 3.5a. Figure 3.5a also shows the true brakeless acceleration of the vehicle, compared to the modeled brakeless acceleration.

The next step (Figure 3.5b) is to subtract the measured acceleration from the modeled acceleration, resulting in an acceleration residual. A Boolean braking state is also required to isolate portions of the series where divergence in measured and modeled acceleration is expected. This Boolean braking status is indicated in Figure 3.5b by gray-shaded regions.

The final step (Figure 3.5c) of the P_{AD}/E_{AD} estimation process of [11] is to obtain the area between the measured acceleration residual (the solid dark blue curve) and the expected acceleration residual if the truck did not brake (the green line). This difference represents the portion of the acceleration of a vehicle that is due to braking, a_{decel}. Referring again to Figure 3.5c, the estimate of a_{decel}, called $\hat{a}_{decel}$, is the length of the green arrows, while the true a_{decel} is the length of the dark blue arrows. Ideally, the green line would lie very close to the dashed dark blue line and have the same mean, because a_{decel} and the area enclosed in Figure 3.5c is linearly related to the power and energy lost to braking:

$$P_{AD} = m_{eff} a_{decel} v \tag{3.10}$$

$$E_{AD} = \int P_{AD} dt \tag{3.11}$$

Once a_{decel} is calculated, the power lost in braking can be calculated. The accuracy of $\hat{P}_{AD}$ is hampered when the acceleration signal is not exact. Acceleration for this method is an estimate based on the wheelspeed of the vehicle [103]. As such, there is noise and delay in the acceleration signal, which, according to Figure 3.5d, has a clear effect on $\hat{P}_{AD}$. The effects of noise and delay in the acceleration signal on the accuracy of $\hat{P}_{AD}$ will be explored in detail later.

In the original analyses, the combination of the drag reduction model proposed by [93] and the estimate for E_{AD} demonstrated the ability to explain more than 80% of the variance in NFC_{true} for the experimental data, but only a limited subset of the available data was used[8]. Therefore, we extended the previous work by examining the data of other trucks, expanding the analysis to a larger dataset. The expansion of the previous work in platooning inference ultimately did not change the resulting inference much, so it will not be covered in full here. See Appendix A for the complete regression analysis.

The regression equation to infer platoon benefits after expansion to a larger dataset is given in Equation 3.12:

$$NFCr = \beta_0 + \beta_1 E_{AD} + \beta_2 DRF_4 + \beta_3(\text{fan power}) + \varepsilon \tag{3.12}$$

$$\hat{NFCr} = -0.014 + 0.066 E_{AD} + -0.53 DRF_4 + 0.0039(\text{fan power}) \tag{3.13}$$

A quick interpretation of the coefficients in Equation 3.13 is provided in Table 3.1. In addition to braking and drag reduction terms ($\hat{\beta}_1$ and $\hat{\beta}_2$), an estimate of the effect of fan power was also included ($\hat{\beta}_3$), as it significantly increased the R^2.

Table 3.1: Interpretation of model coefficients and their 95% confidence intervals.

Coefficient	Units	Value	Effect on Fuel Consumption (NFC_{true})
$\hat{\beta}_1$	$^1/_{\frac{\mathbf{kJ}}{\mathbf{kg \cdot hr}}}$	0.066	1 $\frac{\mathbf{kJ}}{\mathbf{kg \cdot hr}}$ of E_{AD} yields a $6.6 \pm 0.6\%$ increase
$\hat{\beta}_2$	$^{100}/_{\text{drag reduction \%}}$	-0.53	A 10% drag reduction yields a $5.3 \pm 0.5\%$ decrease
$\hat{\beta}_3$	$^1/_{\mathbf{kW}}$	0.0039	10 **kW** of fan power yields a $3.9 \pm 0.7\%$ increase

Even with the expansion of the data, it still only includes a limited range of truck masses, all at a single speed. Furthermore, we see that the model is empirical, which means that there are no guarantees that it will apply to other datasets. To further evaluate the effectiveness of the methods proposed by [10, 11], we will next test the methods on separate unseen data that was collected at a different speed and truck mass.

[8]Specifically, from the truck A2 at the American Center for Mobility (ACM) test track, which is part of the yet-to-be introduced Dataset 1.

3.2.2 Application of Previous Methods to Unseen Data

The selected model will be applied to unexplained poor platoon performance from an Auburn/National Research Council of Canada study [42]. As Table 9 of [42] shows, a 75-foot platoon saved a measured $10.7 \pm 1.2\%$ fuel, but the same 75-foot platoon with regular passenger vehicle cut-ins did not experience measurable savings ($1.0 \pm 1.8\%$). Some important differences exist between this test case and the model source data:

1. The truck was loaded much heavier than the trucks in the source data
2. The test speed was 65 mph instead of the 45 mph speed of the source data
3. The test was conducted at a different track than the source data

It is bold to assume that, in such a different scenario, the model will predict well. After all, aerodynamic drag is a function of velocity squared, which means that aerodynamic drag was over twice as high in the test set. Also, there is only a single run, so this is a proof-of-concept more than a validation.

The inferred effect of E_{AD} on the fuel consumption of the test run will be described in four steps, starting from the raw calculation of E_{AD}, and ending with the prediction of benefits as compared to the published fuel savings.

1. The absolute active deceleration energy must be calculated for each run. Figure 3.6a shows how the cut-in test accumulated braking energy E_{AD} at a steady pace, whereas the other 75-foot runs had no active deceleration.

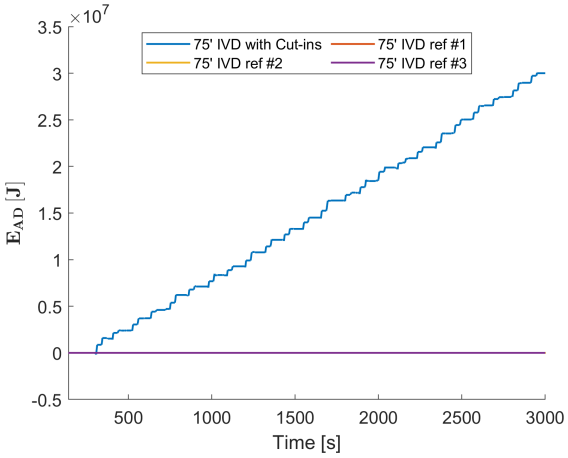

(a) Accumulated raw braking losses in **J** for the data from [42].

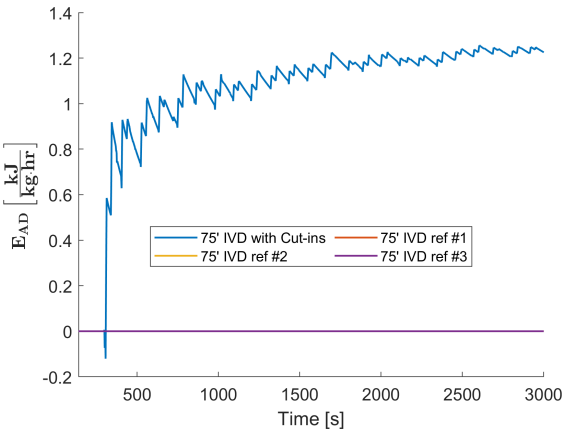

(b) The rolling E_{AD} for the data from [42].

2. Raw E_{AD} in **J** must be converted to mass- and time-normalized E_{AD} with units of $\frac{\mathbf{kJ}}{\mathbf{kg \cdot hr}}$ using the truck mass and elapsed time. Figure 3.6b shows that it takes more than 25 minutes for this to reach a steady value. These values can be input directly into the regression equation.
3. The model is used to calculate the marginal impact of E_{AD} on fuel savings on a rolling basis. Figure 3.7 shows the model prediction for the impact of E_{AD} on fuel savings. The actual measured difference in savings was $9.7 \pm 2.2\%$, which is also plotted. We see that the prediction of the model overlaps significantly with the confidence interval of the difference measured at the end of the test, $9.0 \pm 0.7\%$.

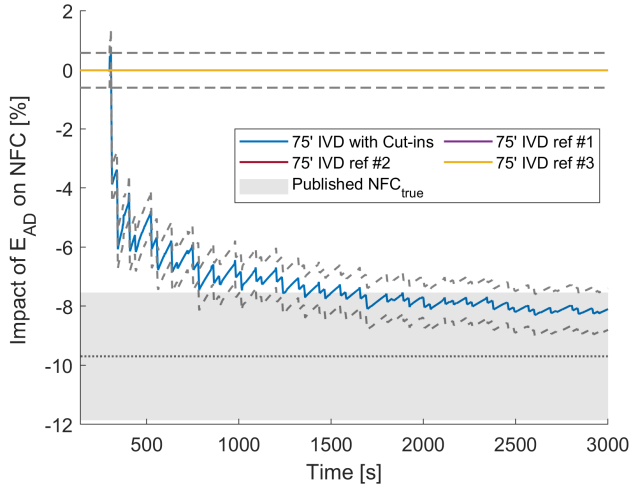

Figure 3.7: The rolling impact of E_{AD} on fuel savings for the cut-in data of [42], using the approach in [11]. The published true marginal impact of braking is shown for comparison. The intervals shown are the 95% confidence intervals.

4. The drag reduction ratio and fan power of the data are used to calculate $NFC_{inferred}$. As Figure 3.8a shows, the previous method does not correctly predict the measured fuel savings. However, it does predict correct trends (the 75-foot platooning saved fuel, whereas the cut-in test did not).

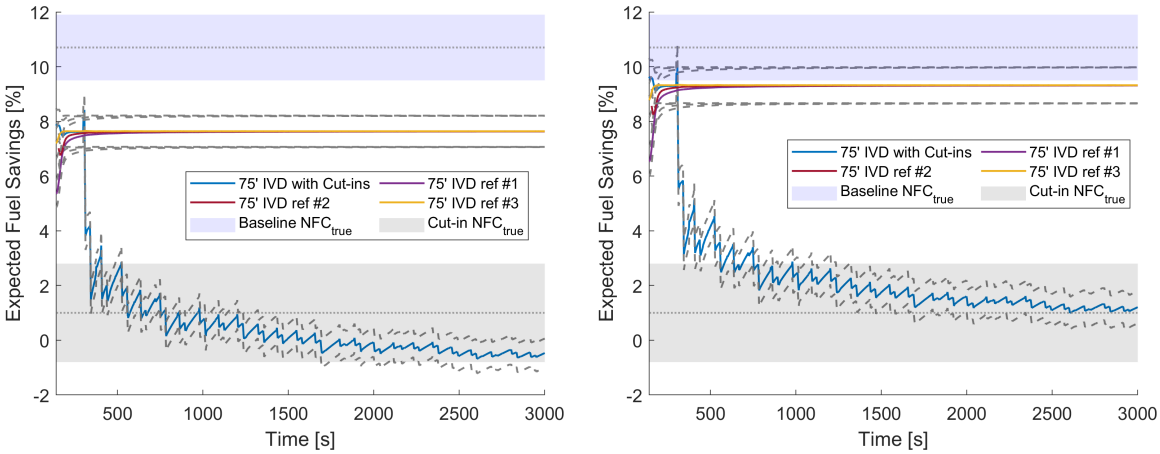

(a) Rolling predicted fuel savings.

(b) Rolling predicted fuel savings with a rough adjustment for proportionality applied.

Figure 3.8: NFC_{true} for the cut-in data in [42] versus $NFC_{inferred}$ calculated by [11].

The failure of the model to correctly predict the measured fuel savings (Figure 3.8a) is likely due to the difference in test speed/truck mass. An incorrect assumption that has been exposed here is that the drag reduction coefficient β_2 is constant. Although the trucks in the source data do have varying masses, the range is narrow and does not overlap with the truck mass of the Canadian trial. Furthermore, the entire dataset was collected at 45 miles per hour. With a 65 mph test case, it appears that $\hat{\beta}_2$ is downward biased: drag was more important here than the model asserts.

By way of troubleshooting, the following approximate proportionalities apply to road vehicles:

1. drag losses $\propto$ speed^2
2. other losses $\propto$ mass, accessory loads nonwithstanding

Knowing that the mass went from 37000 to 65000 lbs, and the speed went from 45 mph to 65 mph, the adjustment to $\hat{\beta}_2$ would be:

$$\frac{65^2}{45^2} \div \frac{65000}{37000} = 1.19, \hat{\beta}_{2,\text{test}} = -0.53 \cdot 1.19 = -0.633 \tag{3.14}$$

Figure 3.8b shows that the model now places the predictions within the range of the measured intervals. When proportionality is applied to the prediction, the model performs better, although it still appears to be conservative. The ad hoc and imprecise nature of the proportionality adjustment leaves something to be desired, but it is rooted in vehicle energy principles.

3.2.3 Discussion

Several things can be taken away from this section. First, the application of the existing model for platooning inference should be approached with caution. More development will be required to apply the methods of [10, 11] to wide-ranging commercial trucking situations

with good accuracy. However, this test case is only a single comparison. Not many statistical conclusions can be drawn from one or two test points.

Despite the downward bias of the model, the model proved illuminating: It correctly predicted whether the platooning truck saved fuel. We can see how useful it would be to infer the platoon energy benefit. Hypothetically, $NFC_{inferred}$ allows platoon operators to immediately know if they have saved fuel from platooning.

The model provides what appears to be a conservative estimate of platoon fuel savings in this case study. However, our aim is to gain a more profound understanding of how braking power and drag reduction affect NPC in a way that is applicable to all possible masses and speeds at which a truck might platoon. To do so, we used a wide sweep of vehicle simulations to query the underlying mechanics and develop an inference framework that is truly physics-based.

3.3 Model Formulation

Based on the adequacy check in the previous section, the relationship found in [10, 11] may not generalize to other speeds and loads of trucks. A more fundamental understanding is required to formulate a model that will apply across all platoons. To unravel the mechanism of platooning benefits, we turn to simulation.

3.3.1 Simulation Setup

The purpose of the simulation was to investigate how the platoon benefit changed according to four variables: mean braking power, drag reduction ratio, speed, and mass. To construct the simulation, a digital twin of the American Cener for Mobility (ACM) Highway loop in Ypsilanti, MI was created in the simulation software IPG TruckMaker. The creation of this virtual test track is documented in Appendix C.

The platooning longitudinal control was accomplished using the built-in IPG Truckmaker ACC, re-parameterized to mimic an aggressive CACC response. The chosen ACC parameters are summarized in Table 3.2.

Table 3.2: The parameters used for the simulated ACC controller.

Parameter	Default value	Used value
Brake Threshold	0.2	0.0
Initial Time Distance [s]	1.8	0.8
Minimal Distance [m]	20	0
Minimal Acceleration [m/s^2]	-2.5	-3.0
Maximal Acceleration [m/s^2]	1.0	3.0
Distance Controller Factor *kd*	36.0	8.0
Distance Controller Factor *kv*	2.0	1.0
Velocity Controller Factor *kv*[1]	13.0	1.0

[1] It is unclear from the product documentation what distinguishes this *kv* from the distance controller *kv*. The best guess is that it represents the acceleration gain (the derivative term of the PID controller).

As for the leader speed controller, the provided IPG Driver was used, which represents a stock cruise control system. By default, the IPG Driver will brake.

In an attempt to cover the full range of drag reduction ratios, platoonable speeds, and truck masses, a full factorial matrix was required, which led to far too many simulations to run manually ($\mathcal{O} \sim 10^4$). Furthermore, the free academic license provided could not simulate both platoon vehicles in parallel, so the leader and follower had to be run sequentially: first, a leader simulation was conducted, then a truck was made to follow behind a traffic object using the leader simulation's velocity trace. Therefore, simulation automation was required. The scripting tools in IPG Truckmaker provided such a solution to automate the simulation. Appendix C also provides the pseudocode that was used to run through the simulation matrix.

The particular settings for the mass, speed, and drag reduction ratio of the simulation are given in Table 3.3. The ranges were selected to cover the domain of platoonable speeds

(45 to 70 mph / 72 to 116 km/h), truck masses (33000 to 77000 lbs / 15000 to 35000 kg), and drag reductions of up to 50%.

Table 3.3: Simulation Mass/Speed/DRR Settings.

Setting	Minimum	Increment	Maximum	Total
Platoon Leader				
Ego Mass [kg]	15,000	by 2000	35,000	11
Desired Speed [km/hr]	72	by 4	116	12
Drag Reduction Ratio [unitless]	-	-	1.0	
Total No. Cases				132
Platoon Follower				
Ego Mass [kg]	15,000	by 2000	35,000	11
Leader Mass [kg]	15,000	by 2000	35,000	11
Desired Speed [km/hr]	72	by 4	116	12
Drag Reduction Ratio [unitless]	0.5	0.1	1.0	6
Total No. Cases				8712
Brakeless Reference				
Ego Mass [kg]	15,000	by 2000	35,000	11
Desired Speed [km/hr]	70	by 1	116	47
Drag Reduction Ratio [unitless]	-	-	1.0	
Total No. Cases				517
Total All Cases				9361

The parameters for the simulated truck and its environment are listed in Table 3.4. Some parameters were not directly specified, so thus had to be back-calculated. Variations in rolling resistance and drag coefficients are due to slight mass/velocity dependence within the tire and drag models.

Table 3.4: Simulation truck and environment parameters.

Parameter	Symbol	Value
Air density	ρ_a	1.205 kg/m^3
Ambient pressure	P_{amb}	1.013 bar
Gravitational acceleration	g	9.806 m/s^2
Frontal area	A_f	10 m^2
Coefficient of drag	C_d	0.843 to 0.848
Coefficient of rolling resistance	C_{rr}	0.00952 to 0.010

The drag and rolling resistance parameters are higher compared to those of a typical modern truck. Furthermore, there are some scenarios simulated that would not ever occur in reality (e.g. a 50% drag reduction at medium-range IVD of 25m / 0.8 s). However, this does not impact the overall generalizability of the simulation. Even if the simulation were re-parameterized to reflect more realistic scenarios, all that this would do is shift the power consumption of the simulated trucks to a different operating regime, which the framework is designed to handle. Because the drag and rolling resistance are higher than usual, the relative impact of drag reduction is somewhat higher than the relative impact of braking than it would be in practice, but this dynamic will be captured by the framework.

The simulations took several days to run on a desktop computer, given the single-node license. At the end, there were 9361 files, each containing 10 Hz data of the simulated truck's performance. However, these files do not contain everything needed to start inferring benefits on the simulated platoons; the NPC_{true} and NFC_{true} of each individual simulation are not a direct output since they are ratios of two separate runs. In the next section, we will detail how NPC_{true} was calculated for each simulation.

3.3.2 Normalized Power Consumption Calculation

The reader will recall that NPC_{true} (or alternatively NFC_{true}) is a ratio of platoon to reference power consumption ($\bar{P}_{plat}$ to $\bar{P}_{ref}$). To obtain NPC_{true} of the simulated trucks, simulations had to be run for both platoons and references. The references were set to have no drag reduction or braking, making them "brakeless" references.

It would be simple to calculate NPC_{true} if the simulation set speed corresponded to the actual mean speed, but this is not the case. The simulated driver was unable to match the desired average speed due to the hills of the virtual environment. This led to a 1-2% difference between the set speed and the actual mean speed, as shown in Figure 3.9.

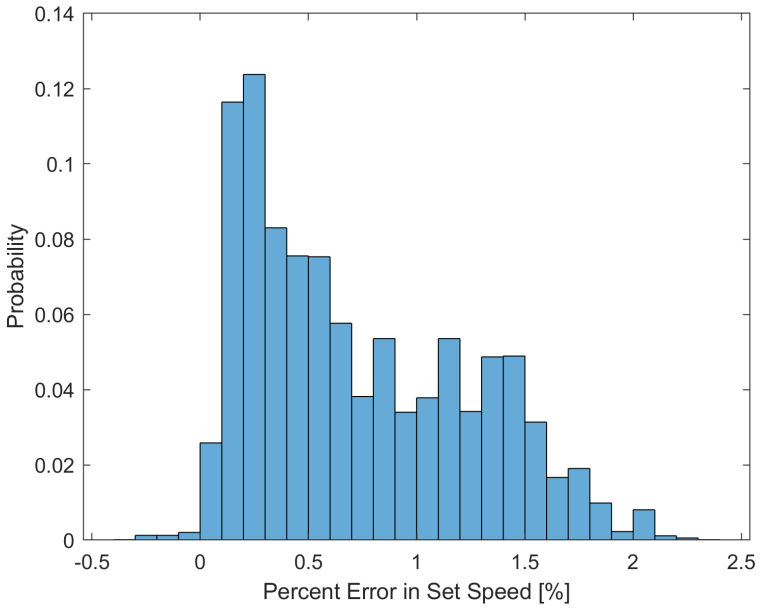

Figure 3.9: Discrepancy between the simulation set speed and the actual mean speed.

Faced with slight mismatches in the set speed and the actual mean speeds, a model was built to predict the reference power consumption for a given mass and mean speed, $\bar{P}_{ref}(m, \bar{v})$. This ensured that the reference had the same mean velocity as that of the platoon. To create the lookup, a two-layer neural network was fit to the brakeless power consumption results. The neural network training was performed automatically using the MATLAB function `fitrnet`, with the selected architecture shown in Table 3.5. The neural net is surely much larger than necessary, but it captured the shape of the reference power

consumption as a function of mass and speed very well. Concerns about overfitting can be overcome by inspecting the actual fit in Figure 3.10, which is smooth and shows no signs of overfitting.

Table 3.5: Reference mean power consumption neural network architecture.

Parameter	Value
Activations	ReLU
Lambda	7.67
Layer Sizes	[299 8]

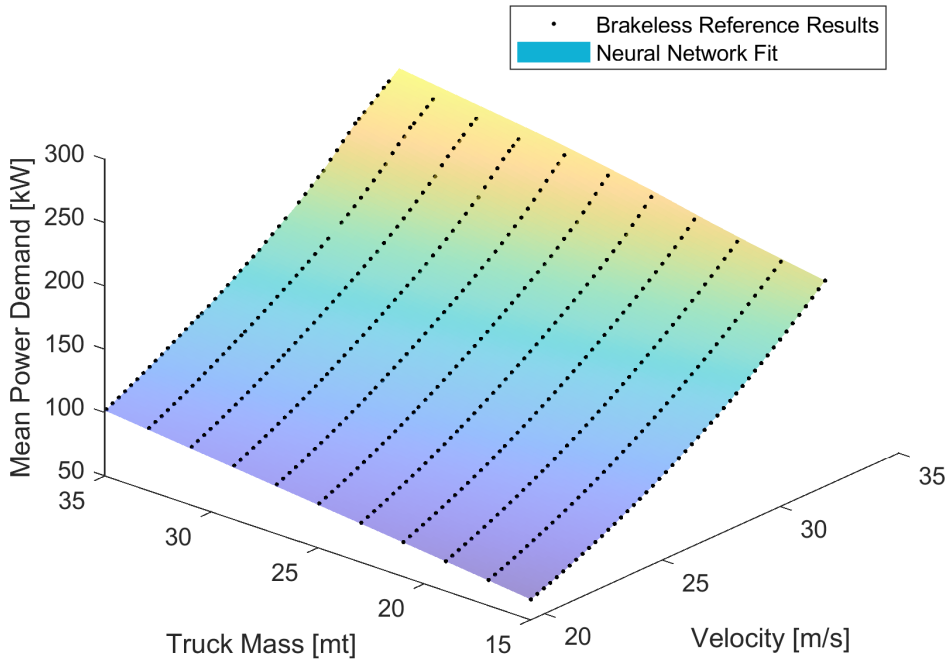

Figure 3.10: Brakeless reference mean power consumption results from simulation and the accompanying neural network model to lookup $\bar{P}_{ref}(m, \bar{v})$.

As seen in the contour plot of Figure 3.11, the error of the regression is less than a kilowatt throughout the range. This means that the brakeless, full-drag power demand of the truck can be predicted throughout the simulated range to within a kilowatt using this equation.

In summary, to obtain the true platooning benefit of the simulations (NPC_{true}), we had to build a lookup of the reference power consumption $\bar{P}_{ref}(m, \bar{v})$. This lookup was generated using a neural network. Ultimately this enables the calculation of the NPC_{true}

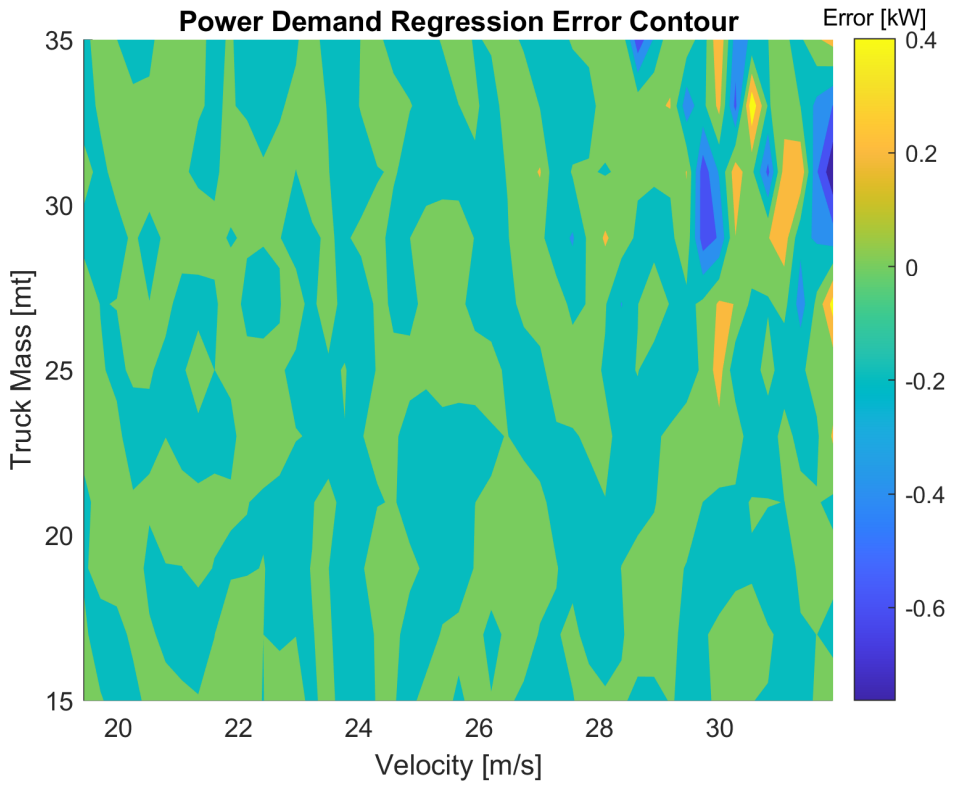

Figure 3.11: Error contour for the neural network fit of $\bar{P}_{brakeless}(m, \bar{v})$.

of any platoon simulation, which we will need to do to validate our proposed inference. Up next, the form of that inference will be developed.

3.3.3 Derivation of the Platooning Energy Consumption Function

In this section, the method for inferring platoon benefits is derived. We start by considering the distribution of the tractive power used to overcome each road load. The overall power of the vehicle (P_{total}) consists of the sum of aerodynamic power (P_{aero}), rolling resistance power (P_{rr}), accessory power ($P_{accessory}$), gravitational power (P_{grvt}), and power related to acceleration/deceleration (P_{accel}/P_{AD}):

$$P_{total} = P_{aero} + P_{rr} + P_{accessory} + P_{grvt} + P_{accel} - P_{AD} \tag{3.15}$$

We can leverage the assumptions made earlier in the chapter reformulate total vehicle power into an expression of platooning benefits. It is assumed that the only difference in energy use between the platoon and the reference vehicle lies in the reduction of braking power and drag, as stated in Assumption 1. Furthermore, the mean power of each term in Equation 3.15 reflects the mean total power $\bar{P}_{total}$ (based on Assumption 5). Next, any deceleration power (P_{AD}) will be matched with a later acceleration event P_{accel} during the microtrip (per Assumption 3)[9]. Therefore, NPC may be expressed as:

$$NPC = \frac{DRR \times \bar{P}_{aero} + \bar{P}_{AD} + \bar{P}_{rr} + \bar{P}_{accessory} + \bar{P}_{grvt}}{\bar{P}_{aero} + \bar{P}_{rr} + \bar{P}_{accessory} + \bar{P}_{grvt}} \tag{3.16a}$$

Decomposing Equation 3.16a into two terms, one for drag reduction, and one for braking losses, yields Equation 3.16b:

$$NPC = \frac{DRR \times \bar{P}_{aero} + \bar{P}_{rr} + \bar{P}_{accessory} + \bar{P}_{grvt}}{\bar{P}_{aero} + \bar{P}_{rr} + \bar{P}_{accessory} + \bar{P}_{grvt}} + \frac{\bar{P}_{AD}}{\bar{P}_{aero} + \bar{P}_{rr} + \bar{P}_{accessory} + \bar{P}_{grvt}} \tag{3.16b}$$

[9]Theoretically, one could find the positive analogue of P_{AD}, although it would be difficult to isolate the portion of positive acceleration that was due to braking from the rest of it.

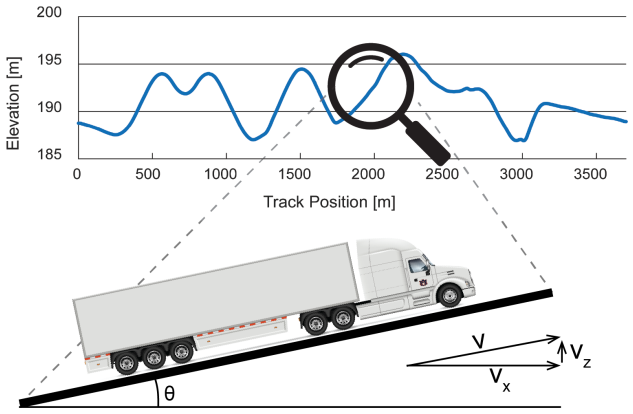

Figure 3.12: Components of velocity over grade. Elevation is not to scale.

In fully expanded, uncanceled form, Equation 3.16 is:

$$NPC = \frac{\int \left[DRR \times 0.5\rho_a C_d A_f (v+v_w)^2 + m_{eff} a_{decel}\right] \cdot v + C_{rr} m g v_x + P_{accessory} + m g v_z \; dt}{\int \left[0.5\rho_a C_d A_f (v+v_w)^2\right] \cdot v + C_{rr} m g v_x + P_{accessory} + m g v_z \; dt} \tag{3.16c}$$

Within Equation 3.16c, v_w is the headwind wind velocity in the vehicle frame, a value that is usually not available. Grade angles are what distinguishes v_x and v_z from v, as the instantaneous relationship between the three is Pythagorean, as shown in Figure 3.12. Over the course of a microtrip, the mean gravitational power is proportional to the average vertical velocity, $\bar{v}_z$, and the mean rolling resistance is related to the average horizontal velocity in the earth-fixed frame, $\bar{v}_x$.

For the drag reduction part of Equation 3.16b (the first term), it is illuminating to see how the term will change across masses and speeds. To do so, the drag fraction of power can be examined, which is defined as the fraction of power that is spent to overcome aerodynamic drag:

$$\textbf{drag fraction} \triangleq \left(\frac{\bar{P}_{aero}}{\bar{P}_{aero} + \bar{P}_{rr} + \bar{P}_{accessory} + \bar{P}_{grvt}} \mid DRR = 1.0 \right) \tag{3.17}$$

Consequently, the term belonging to drag reduction in equation 3.16b, which we will call the Power Reduction Ratio (PRR), can be expressed in terms of **drag fraction**:

$$PRR \triangleq \frac{DRR \times \bar{P}_{aero} + \bar{P}_{rr} + \bar{P}_{accessory} + \bar{P}_{grvt}}{\bar{P}_{aero} + \bar{P}_{rr} + \bar{P}_{accessory} + \bar{P}_{grvt}} = \textbf{drag fraction} \times (DRR - 1) + 1 \quad (3.18)$$

In Figure 3.13 **drag fraction** has been plotted versus speed and mass, presuming accessory and gravity loads are zero[10].

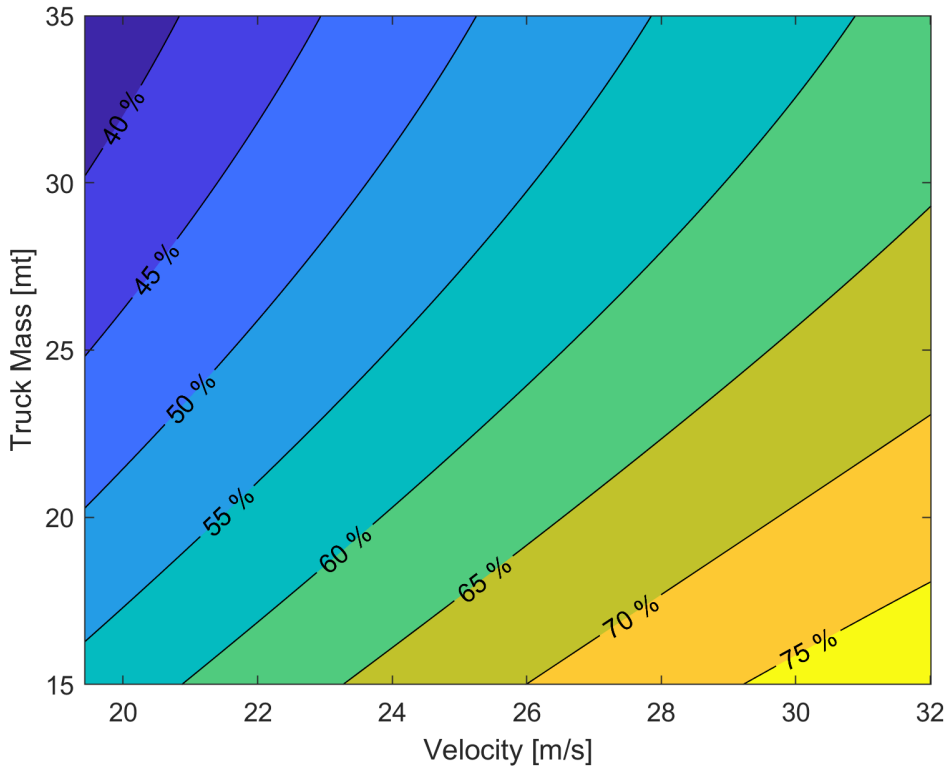

Figure 3.13: The **drag fraction** as a function of mass and mean speed, with no drag reduction (DRR=1.0). Values from simulation.

The important takeaway from Figure 3.13 is that drag reduction's efficacy changes with speed and mass. According to Figure 3.13, an unloaded truck at 20 m/s (45 mph) has

[10]To derive a continuous **drag fraction** from the results at discrete simulated masses and velocities, a second-order linear regression was fit:

$$\textbf{drag fraction} = \beta_0 + \beta_1 \bar{v} + \beta_2 \bar{v}^2 + \beta_3 \bar{v}\bar{m} + \beta_4 \bar{m} + \beta_5 \bar{m}^2 \quad (3.19)$$

The regression of Equation 3.19 yielded an R^2 of 1.0, meaning it can be confidently used to look up the drag fraction at a given velocity and mass.

approximately the same drag fraction as a loaded truck at 30 m/s (65 mph). With the addition of loads such as air conditioning, the percentages shown in Figure 3.13 would be scaled down proportionally according to Equation 3.17. In contrast to the Power Reduction Ratio, which is the first term of Equation 3.16b, the second term may be defined as the Power Increase Fraction:

$$PIF \triangleq \frac{\bar{P}_{AD}}{(\bar{P}_{aero} + \bar{P}_{rr} + \bar{P}_{accessory} + \bar{P}_{grvt})} \tag{3.20}$$

Although all of the preceding equations starting with Equation 3.16 could be used to infer platooning energy savings, by substituting the total power consumption of the platoon vehicle into Equation 3.16, a more practical expression is found. To disambiguate the origin of each term, the subscripts "plat" and "ref" will be applied to each term.

Total platooning engine power is the numerator of Equation 3.16. The total engine power infers many terms of Equation 3.16, leading to Inferred NPC:

$$NPC_{inferred} \triangleq \frac{\bar{P}_{total,\,plat}}{\bar{P}_{total,\,plat} + (\bar{P}_{aero,\,ref}(1 - DRR) - \bar{P}_{AD,\,plat})} \tag{3.21a}$$

Returning to the discussion of what constitutes a reference, a realistic reference may be a better choice than a reference without brakes. Suppose that the realistic reference mean braking power is known, which we shall call $P_{AD,\,ref}$. Since the realistic platooning benefit depends on the difference between platoon braking and reference braking, $\bar{P}_{AD,\,ref}$ should oppose P_{AD} in the denominator, like so:

$$NPC_{inferred} = \frac{\bar{P}_{total,\,plat}}{\bar{P}_{total,\,plat} + (\bar{P}_{aero,\,ref}(1 - DRR) - \bar{P}_{AD,\,plat} + \bar{P}_{AD,\,ref})} \tag{3.21b}$$

$$\bar{P}_{AD,\,ref} \triangleq m_{eff,\,plat} \cdot a_{decel,\,ref} \cdot v_{ref} \tag{3.21c}$$

Note that $\bar{P}_{AD,\,ref}$ uses acceleration and velocity of the reference vehicle, but the mass of the ego vehicle. In this way, the braking loss characteristics of the reference vehicle with a

different mass may be applied to the ego vehicle. For example, a heavier lead vehicle could be used as a reference for the lighter follow vehicles in a platoon.

One nice property of Equation 3.21 (in all its forms) is that it implicitly includes the change in kinetic and potential energy. Although $NPC_{inferred}$ is algebraically identical to Equation 3.16, it is less dependent on the deterministic models of an imaginary truck because $\bar{P}_{total,platoon}$ is a measurable quantity.

Although power consumption is of interest, fuel consumption is often the most important figure of merit. Therefore, $NFC_{inferred}$ can be expressed in a similar form $NPC_{inferred}$, although $\bar{P}_{AD}$ and $\bar{P}_{aero}$ will need to be expressed in terms of fuel. A fuel conversion factor κ is defined to convert wheel power into fuel power energy:

$$\kappa \triangleq \frac{1}{\eta_{gen,\,plat} LHV \rho_{fuel}} \tag{3.22}$$

with fuel conversion efficiency $\eta_{gen,\,plat}$, Lower Heating Value LHV, and fuel density ρ_f. Furthermore, the efficiency of the powertrain may be altered by platooning, so a multiplier α is defined as the ratio of $\eta_{gen,\,ref}$ to $\eta_{gen,\,plat}$

$$\alpha \triangleq \frac{\eta_{gen,\,ref}}{\eta_{gen,\,plat}} \tag{3.23}$$

Using κ, α, and fuel rate F allows the conversion of $NPC_{inferred}$ to $NFC_{inferred}$:

$$NFC_{inferred} \triangleq \alpha \frac{\bar{F}_{plat}}{\bar{F}_{plat} + \kappa(\bar{P}_{aero}(1 - DRR) - \bar{P}_{AD} + \bar{P}_{AD,\,ref})} \tag{3.24}$$

According to Assumption 2, the difference in the efficiency of the powertrain by platooning is neglected, leaving $\alpha = 1$. Although inclusion of such information is valuable, the characteristics of powertrain efficiency depend on architecture, and this study lacked sufficient information to make such a judgement about the appropriate value of α.

Finally, to include the effect of regenerative braking in braking losses, a small adjustment is necessary. $\bar{P}_{AD}$ need only be multiplied by the fraction of braking energy lost to heat $\eta_{AD,\,lost}$, which is the complement of the regenerative fraction $\eta_{AD,\,regen}$.

$$\bar{P}_{AD,\,lost} = \eta_{AD,\,lost}\bar{P}_{AD} = (1 - \eta_{AD,\,regen})\bar{P}_{AD} \tag{3.25}$$

Regenerative braking vehicles were not analyzed here, so $\eta_{AD,\,lost}$ was always one, and Equation 3.25 was not necessary.

Ratios can be difficult to use in practice, as their statistics are often dubious and misleading compared with those of differences in means and other common statistics for comparing two groups. If a relative energy benefit is not strictly required, then $NPC_{inferred}$ and $NFC_{inferred}$ can be expressed as an absolute difference instead, which will still give the ability to determine whether fuel savings were achieved. This difference, $\Delta\bar{P}_{inferred}$, is easily extracted from Equation 3.21b.

$$\Delta\bar{P}_{inferred} \triangleq \bar{P}_{aero,\,ref}(1 - DRR) - \bar{P}_{AD,\,plat} + \bar{P}_{AD,\,ref} \tag{3.26}$$

Multiplying $\Delta\bar{P}_{inferred}$ by κ gives the fuel equivalent, assuming $\alpha = 1$. In the results chapter, $\Delta\bar{P}_{inferred}$ may be used at times in place of $NPC_{inferred}$ when the ratio statistics are difficult to construct.

This concludes the derivation of $NPC_{inferred}$ and $NFC_{inferred}$ as functions of mass, velocity, drag reduction, and braking losses. The simulation results will now be used to demonstrate the validity of the relationships derived here.

3.3.4 Verifying the Platooning Energy Consumption Function with Simulation

Equation 3.16 and its alternative forms have been claimed to describe the energy consumption of the platoon relative to a non-platooning reference. In this section, NPC_{true} from the simulation will be compared to $NPC_{inferred}$ from Equation 3.16c. Recall that NPC_{true} and $NPC_{inferred}$ are two different things here: NPC_{true} is derived from the ratio of net total engine power of two different simulations, whereas $NPC_{inferred}$ can use one simulation. Of the two, only $NPC_{inferred}$ can be calculated in application conducted test replicates.

To begin, we shall look at the results of the entire simulation space in raw form. The mean power consumption $\bar{P}$ of all simulation data is plotted in Figure 3.14. The entire range from Table 3.3 is included in Figure 3.14, which means that the drag reduction to 50% is shown. At a given mass and speed, there are often points above and below the brakeless

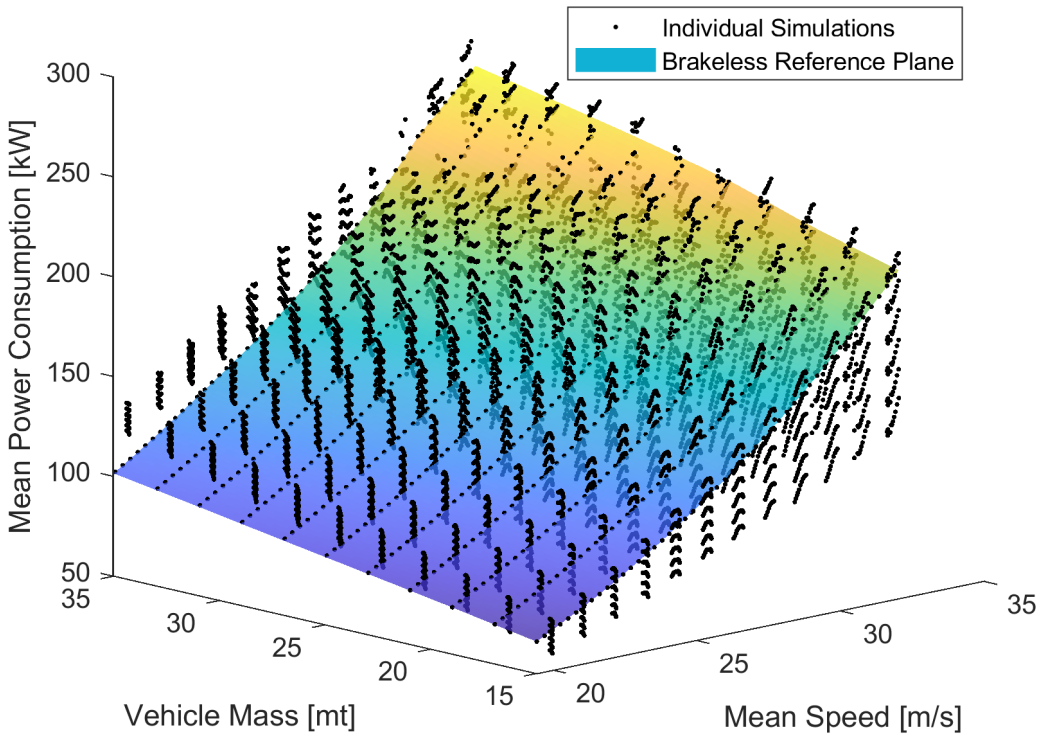

Figure 3.14: Mean power consumed for all simulation runs, including the brakeless reference runs.

reference plane, which will lead to NPC above and below 1.0. Vertical variation in points comes from the combination of other simulation inputs: the leader mass, the drag reduction ratio, and the amount of braking. About 6% of the 9361 simulations had a large difference in starting and ending speeds, making them ill-suited for comparison. Consequently, these simulations were removed using the MATLAB function `findoutliers`, which uses 3 median absolute deviations (MAD) to detect outliers[11]. After this outlier removal, the difference in starting and ending velocity for the remaining 8831 simulations was less than 0.5 m/s absolute, 1.5% relative. The sparsity in points in the 30 m/s, 30 mt region near the top of Figure 3.14 is due to this outlier removal.

Figure 3.15 illustrates the fluctuations of NPC_{true} versus mass and speed. In the left

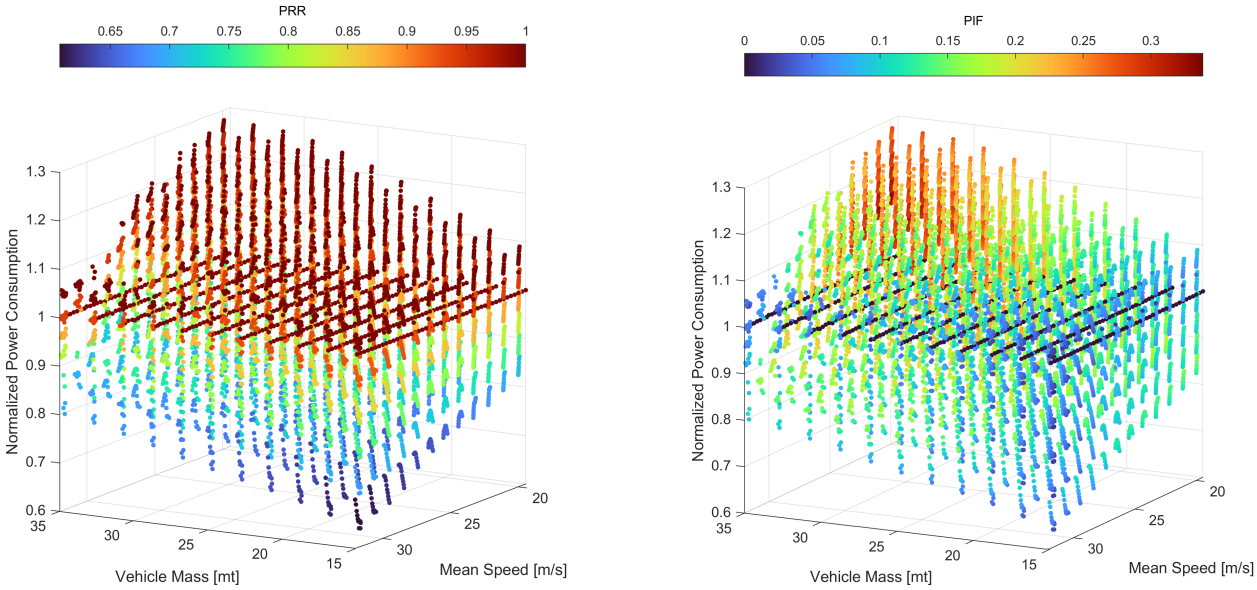

(a) Colored by Power Reduction Ratio due to drag reduction, Equation 3.18.

(b) Colored by Power Increase Fraction due to braking, Equation 3.20.

Figure 3.15: Normalized Power Consumption versus mass and speed.

panel, the data points are color-coded based on the Power Reduction Ratio (PRR) resulting from drag reduction, while in the right panel, they are color-coded based on the Power Increase Fraction (PIF) due to braking. Recall that the metric of interest $NPC_{inferred}$ is simply the sum of PRR and PIF, which are colored in Figure 3.15.

[11]The outlier removal was applied to the percent difference in initial and final speed, $(v_f - v_i)/\bar{v}$.

Finally, the NPC_{true} can be compared to $NPC_{inferred}$ to verify that the trend holds. Figure 3.16 shows that the trend is as expected, with the data falling closely along a 1:1 line. The R^2 of the line in Figure 3.16 is 0.998. The small deviations off the line in Figure 3.16 are likely due to some trivial loads, such as chassis rebound losses and slight inaccuracies in the neural network power lookup for $\bar{P}_{ref}$.

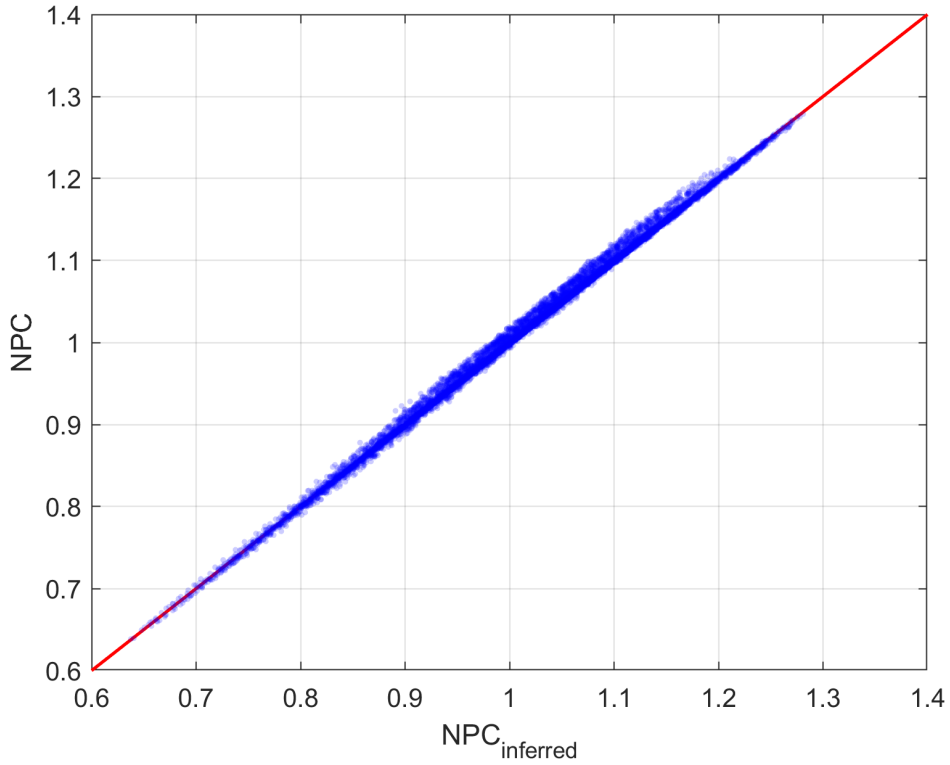

Figure 3.16: NPC_{true} versus $NPC_{inferred}$ for the simulation. $R^2 = 0.998$.

Based on these results, there is a high level of confidence that $NPC_{inferred}$ can generate a good inference of NPC_{true}. However, the methods developed on the basis of this simulation study have yet to be proven in application, beyond the previous work of the author with E_{AD}. In particular:

- In practice, can $NPC_{inferred}$ be calculated?
- If so, how good do the estimates of DRR and P_{AD} need to be to provide useful information? For example, is the method good enough to tell whether platooning saved any energy?

The following section will concentrate on the practical approximation of $NPC_{inferred}$. After practical ways to generate platoon energy inferences in real-time have been established, the chapter concludes with a description of the experimental datasets and the data analysis procedure, in preparation for the results chapter.

3.4 Practical Calculation of Platooning Energy Change

If $NPC_{inferred}$ or $NFC_{inferred}$ can be reliably calculated using vehicle signals, then in theory the effect of platooning on energy consumption can be calculated. As stated above, Equation 3.21 is preferred to Equation 3.16. Equation 3.16 leaves all terms except DRR and $\bar{P}_{AD}$ to parameterization, mainly relying on a model of an imaginary vehicle. Only DRR and $\bar{P}_{AD}$ are observed if Equation 3.16 is used directly. On the other hand, Equation 3.21 is based on $P_{total,\,plat}$, which can make use of vehicle ECU estimates and reduce the parameterization error. $P_{total,\,plat}$ is calculated from SAE J1939 CAN signals in the following manner, converting everything into SI units:

1. Use Actual Engine Percent Torque and Reference Torque to get the calculated engine torque at the current speed (SPN 513, 4154, and 544)
2. Use Engine Speed times the Torque to get engine output power (SPN 190)

The engine speed signal is well known, but the percent torque values are known to be inaccurate, especially at low loads [104]. SPNs 513 and 4154 are supposed to include engine friction and thus be representative of engine shaft output, but it is unknown how each manufacturer accounts for this and other dynamics.

It will also be useful to examine $NFC_{inferred}$ in addition to $NPC_{inferred}$, since $NFC_{inferred}$ is based on a different measurement than $NPC_{inferred}$ (SPN 183 or 1600). Like the torque signal, the fuel rate is also prone to errors at low engine loads [105], and the transient accuracy of the CAN fuel rate is questionable.

To proceed, we must calculate $\bar{P}_{aero}$, DRR, and $\bar{P}_{AD}$. Next, we describe the suggested course for calculating these variables from real data.

3.4.1 Aerodynamic Drag Reduction Calculation

The effect of drag reduction on $NPC_{inferred}$ and $NFC_{inferred}$ depends on two pieces of information:

1. the reference aerodynamic power, $\bar{P}_{aero}$
2. the drag reduction ratio DRR

What are the implications if $\bar{P}_{aero}$ or DRR are misspecified? If $\bar{P}_{aero}$ is incorrect, then the reduction in power due to the reduction in drag will be incorrectly scaled in proportion to $DRR - 1$. Therefore, an error in **drag fraction** is of little concern if the amount of drag reduction is small $DRR \approx 1$. However, the error in the reduction in drag power is directly proportional to the error in DRR, which emphasizes the importance of a good DRR model. Ultimately, an incorrect $\bar{P}_{aero}$ and DRR will move the zero-benefit line, but it will not change the marginal effect of braking losses. It is neither within the scope nor the ability of this dissertation to validate the estimates for $\bar{P}_{aero}$ and DRR, which would require high-fidelity aerodynamic simulations and much better truck instrumentation, as in [33]. It is within the scope of this dissertation to show how $\bar{P}_{aero}$ and DRR were calculated here. $\bar{P}_{aero}$ is will be calculated using a rough approximation for headwinds equal to

$$\bar{P}_{aero} \triangleq \frac{1}{\Delta t} \int \frac{1}{2} \rho_a C_d A_f (v + v_w)^2 v \; dt \tag{3.27}$$

The velocity v in Equation 3.27 will simply be the ground velocity. v_w is the headwind velocity in the vehicle longitudinal axis. C_d and A_f are parameters that will be specified for the vehicle based on the aerodynamics of the truck and trailer.

Equation 3.27 uses a very rough approximation for headwind that can be improved upon in future work. Application of this approximation may be noted in [106] and other similar works. With it, the effect of the lateral component of crosswinds is neglected, but in reality crosswinds do cause an increase in vehicle drag [82].

Regarding v_w and air density ρ_a, a tool was developed to query terrestrial winds from the OpenWeatherMap (OWM) API to collect current environmental conditions. The tool will be briefly demonstrated, though the number of allowed API calls limited the ability to use weather queries.

1. With a provided latitude, longitude, Unix time, and API key, the OWM "One Call API 3.0" interface returns a weather data structure at the given time and location, including wind speed and direction and ambient temperature, pressure, and humidity.
2. The queried wind data is estimated at 10 m above the ground, which is far above the mean height of HDVs (2 m). The wind velocity can be translated from 10m to 2m using the log wind profile equation:

$$\frac{v_{w,\,2m}}{v_{w,\,10m}} = \frac{\ln z_{2m}/z_0}{\ln z_{10m}/z_0} \tag{3.28}$$

   z_0 is called the surface roughness parameter, and it represents the ground-plane displacement where wind is assumed to have zero velocity. For highways, [107] gives a value of $z_0 = 0.5$ meters, and for forests, $z_0 = 1.28$ meters. The resulting wind speed multipliers for highways and forests are:

$$v_{w,\,2m} = \frac{\ln z_{2m}/0.5}{\ln z_{10m}/0.5} v_{w,\,10m} = 0.463 \cdot v_{w,\,10m} \tag{3.29}$$

$$v_{w,\,2m} = v_{w,\,10m} \frac{\ln z_{2m}/1.28}{\ln z_{10m}/1.28} = 0.217 \cdot v_{w,\,10m} \tag{3.30}$$

   This suggests that the wind speed at a height of 10 meters above the ground is lower than half, possibly only a fifth, of what a truck encounters. The difference in the values of z_0 stated in [107] highlights a notable contrast between the multiplier for different highway environments.
3. With the wind speed and direction acquired, the next step is to obtain the heading of the vehicle. This is possible through GPS odometry.
4. The vector addition of wind and ground velocity will allow the wind direction to be expressed in the vehicle body frame, instead of the Earth-fixed frame of the queried

data. Vehicle and wind velocity and are expressed in the Cartesian East-North directions using their direction ϕ:

$$v_{E,gnd} = v\sin(\phi) \tag{3.31}$$

$$v_{N,gnd} = v\cos(\phi) \tag{3.32}$$

$$\vec{v}_{apparent} = \langle v_{E,gnd} + v_{E,wind}, v_{N,gnd} + v_{N,wind} \rangle \tag{3.33}$$

$$\phi_{apparent} = \arctan(\vec{v}_{apparent}) \tag{3.34}$$

$$\psi = \phi_{apparent} - \phi_{gnd} \tag{3.35}$$

$$\|v_{apparent}\| = \sqrt{v_{E,apparent}^2 + v_{N,apparent}^2} \tag{3.36}$$

$$\vec{v}_{body\ axis} = \langle \|v_{apparent}\|\cos(\psi), \|v_{apparent}\|\sin(\psi) \rangle \tag{3.37}$$

$$v_w = \|v_{apparent}\|\cos(\psi) - v \tag{3.38}$$

5. Assuming dry air, the air density can be gathered from the ambient temperature and pressure using the Ideal Gas Law. If desired, relative humidity can be used to obtain the density of moist air, though corrections are usually less than 5%.

To demonstrate the OWM method, minute-by-minute weather data was queried for a 3-hour platooning trip from Auburn, AL to Mobile, AL. Despite the minute-by-minute queries, the data does not update that frequently. The GPS trace of the trip is on the map in Figure 3.17a. It is difficult to say what the proper surface roughness z_0 is for the trip, so the resulting longitudinal windspeed v_w and the yaw angle ψ have been plotted for both $z_0 = 0.5$ m and $z_0 = 1.28$ m.

This weather query tool is designed to calculate approximate aerodynamic conditions while platooning. After all, the yaw angle and the wind speed both significantly affect the platooning benefits. To neglect them is necessarily to accept the possibility of a large error in the platoon energy consumption estimates. Other alternatives for including weather effects include physical measurements of aerodynamic conditions (such as underhood anemometers).

(a) GPS plot of a trip from Auburn, AL to Mobile, AL.

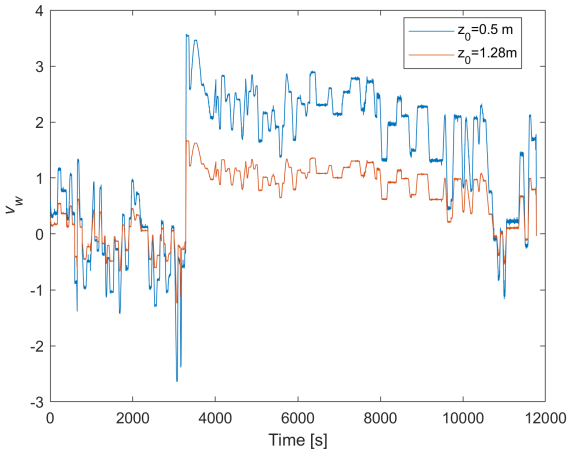

(b) Estimated headwind v_w in the vehicle frame.

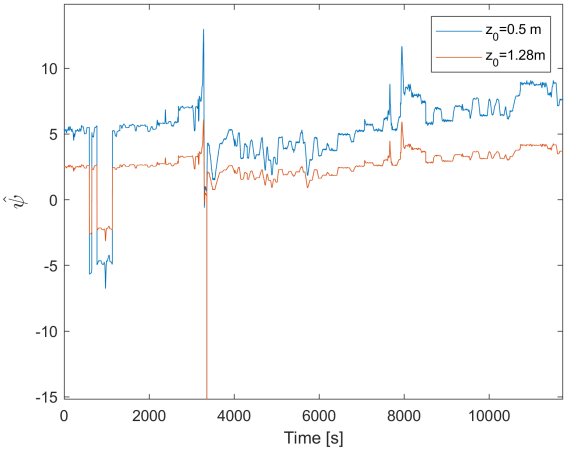

(c) Estimated yaw angle $\hat{\psi}$ in the vehicle frame.

Figure 3.17: A demonstration of using the OpenWeatherMap API to query wind speed and direction on a 3-hour trip from Auburn to Mobile and estimate vehicle wind speed and yaw angle at $z_0 = 0.5m$ and $z_0 = 1.28m$.

3.4.2 Vehicle Acceleration Calculation

Reliable vehicle acceleration signals are very important to the success of the methodology. Estimation of braking power completely depends on the acceleration signal. Should the acceleration signal be clean enough, the vehicle parameters could be adaptively refined. Recent work has fused IMU, GPS, and CAN odometry with asynchronous sample rates into filtered estimates of vehicle longitudinal dynamics [108, 109]. Implementing such a filter was beyond the scope of this work but may improve the acceleration estimates.

Wheelspeed signals, while vulnerable to slight biases and quantization errors, are very reliable. In theory, differentiation of a wheelspeed signal can provide a good acceleration signal. However, signal noise is considerably amplified by differentiation. To use the wheelspeed signal to provide a reliable estimate of acceleration, a filter may be implemented. In previous work, a constant-jerk Kalman filter from [103] was implemented, with good results[12]. An effort was undertaken to integrate the IMU accelerometer data by expanding upon the method proposed by [103], given that all the trucks in this research were fitted with Memsense 3020 accelerometers positioned in the cab. Yet, the inclusion of accelerometer data proved to be ineffective due to the high levels of noise caused by the vibrations of semi trucks. Additionally, not all datasets recorded the accelerometer data, despite being equipped.

Although the Kalman filter of [103] has good smoothing and easy adjustment, it is subject to phase distortion. There is a trade-off between smoothing and delay, as Figure 3.18 shows. As the Q/R ratio is increased, the filter captures transients better and also allows more noise to pass through. Adding to the complexity, the majority of filters exhibit varying phase delays across different input frequencies, as illustrated in Figure 3.18b.

[12]For more details on Kalman filters see [110, 111]

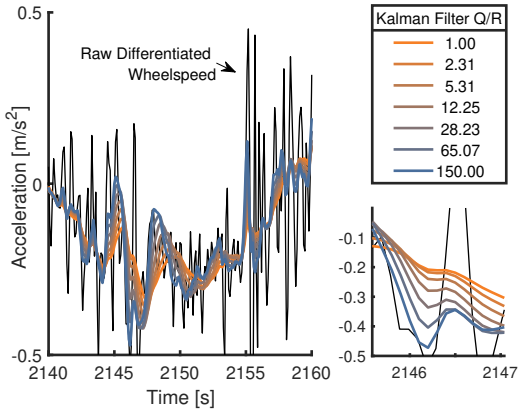

(a) The Kalman filter from [103] that was used in [10, 11] to get an experimental acceleration estimate.

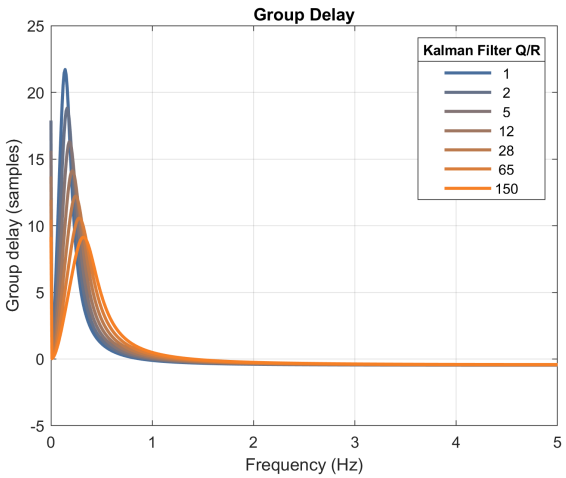

(b) The group delay in samples versus transient frequency of the Kalman filter from [103] across a sweep of Q/R ratios.

Figure 3.18: Illustration of the attentuation and phase distortion of the wheelspeed Kalman filter from [103].

The phase distortion induced by the Kalman filter presents some problems for the task at hand, namely, estimating the braking deceleration. While the acceleration estimates are much less noisy, phase distortion dynamically affects the acceleration residuals. With only a boolean indication of braking losses (like what we have here), this leads to windowing issues: the braking boolean will indicate there is braking, but the acceleration estimate lags behind the boolean.

The options to reduce the effect of phase distortion are to increase sample speed or to use a linear phase filter, such as a Finite Impulse Response (FIR) filter. A linear phase filter is one which has a constant delay, as opposed to the non-constant delay of the Kalman filter in Figure 3.18. For more details on FIR filters see [112]. Increased sample speed is not a possibility here since the experiments are already done. A linear phase filter, on the other hand, can help, because there is only a constant delay, not phase distortion. In other words, the delay induced by a discrete linear phase filter is a constant number of samples, not the frequency-varying curve of Figure 3.18b. Transforming a linear phase filter with a constant delay into a zero-delay filter can be achieved by operating the filter in a retroactive manner,

acknowledging that the updates based on measurements are trailing by the amount of the filter delay. A slight delay in computations (e.g., a second) is deemed permissible within the scope of this study and aligns with the essence of "real-time" processing.

With 10 Hz wheel speed data, Nyquist's sampling theorem states that the fastest transients that can be observed are at 5 Hz. The question is: What frequencies can be cut off safely by the acceleration filter? Longitudinal vehicle dynamics are relatively slow, usually in the 1 Hz range. The order of the filter is also important, as a higher order may respond faster, but at a computational cost. Ultimately, the desired filter should have a low RMSE in terms of the true velocity and estimated velocity, as well as the true acceleration and estimated acceleration.

FIR Acceleration Filter Design

The most important criterion for the filter is that it removes the effect of wheelspeed noise without dampening braking transients, which would downward-bias $\hat{P}_{AD}$. Both a good estimate of acceleration and a consistent estimate of velocity are desired, so filtering will be applied to the differentiated signal. One issue with signal differentiation is that it applies a high-pass filter to the signal. Thus, differentiating a signal and then applying a low-pass filter is somewhat contradictory.

To find good settings for the FIR filter, a grid search was conducted. First, a representative simulation with braking was pulled out for analysis. White noise with a standard deviation of 0.02 m/s was added to the velocity of the simulation vehicle, which is according to the noise statistics derived in the previous work of the author [11]. The wheelspeed data was then numerically differentiated. Then, FIR filters with different cutoff frequencies and filter orders were applied to the differentiated noisy wheelspeed.

Numerical differentiation places the samples half a sample later and reduces the number of samples by one. To account for this, only odd-order FIR filters were considered, which, with their delay of $n/2$ samples also being half a sample late, can realign the samples into

the original time. To maintain realistic applicability, the MATLAB function `filter` was used instead of the `filtfilt`, and the data was shifted into the past and the null spaces discarded. This preserved the interpretation that the filter was being applied to data from $n/2$ samples ago, where n is filter order. The command `filtfilt` runs a complete forward and backward pass and is not precisely equivalent to a shifted output from `filter`.

Filter orders from 1 to 49 and cutoff frequencies from 0.025 to 2.5 were tried. The metrics of interest included RMSE and maximum error in velocity:

$$RMSE_v \triangleq \sqrt{\frac{\sum_{i=1}^{n}(v_{truth,i} - v_{filter,i})^2}{n}} \tag{3.39}$$

$$RMSE_a \triangleq \sqrt{\frac{\sum_{i=1}^{n}(a_{truth,i} - a_{filter,i})^2}{n}} \tag{3.40}$$

$$\text{Maximum Velocity Error} \triangleq \max|v_{truth} - v_{filter}| \tag{3.41}$$

v_{filter}, in this case, is trapezoidally integrated velocity, offset to match the mean truth velocity. A low RMSE indicates that the noise was effectively removed, where a low max absolute error indicates that the velocity transients are not being damped.

The final FIR filter design that attempted to balance noise attenuation with transient fidelity was a 27th-order filter with a cutoff frequency of 1.2 Hz. These parameters struck a good compromise between a low maximum velocity error and a low $RMSE_v$. The filter lags by 1.4 seconds, which is a relatively long delay, but is acceptable for online braking analysis. An alternative design that is more suitable for control tasks may be selected using Figure 3.19, which shows max velocity error in color, and the $RMSE_v$ as a contour. For a filter with less delay but similar performance, an 11th-order filter with a 1.375 Hz cutoff frequency would have similar maximum error and a lag of 6 samples, or 0.6 seconds, which is not so unreasonable that it could not be used for control optimization purposes.

It should be noted that the methods employed here are aimed at obtaining a good acceleration signal for further calculation of $\hat{P}_{AD}$ and possible refinement of model parameters. They are intended to be predictable, which is why an FIR filter was selected over something

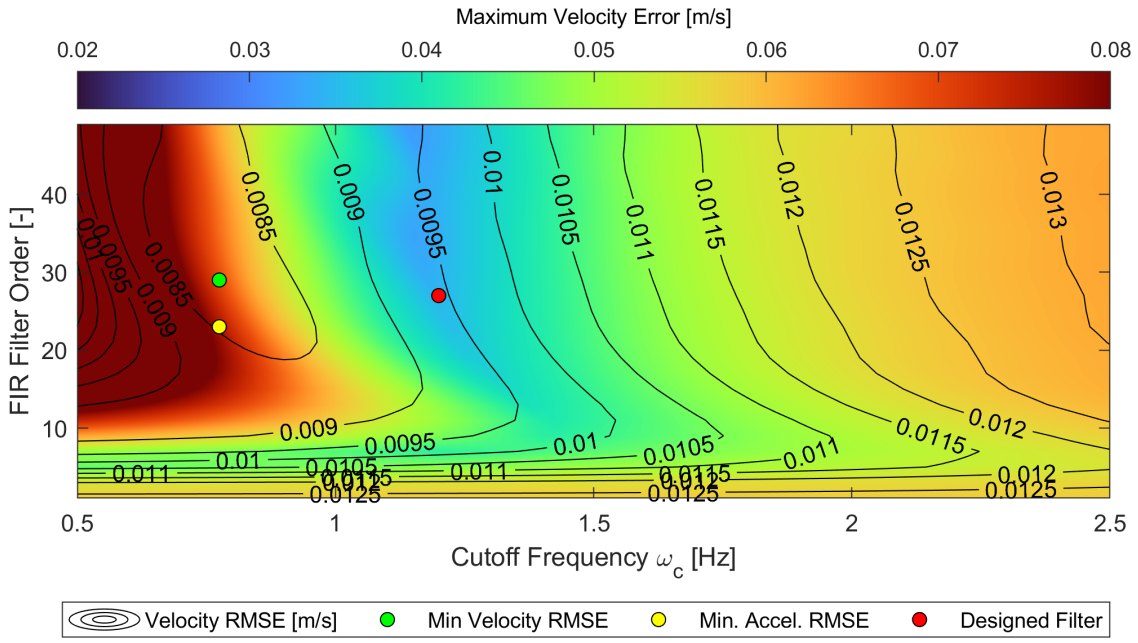

Figure 3.19: FIR filter errors, showing maximum velocity error in color, and the velocity RMSE as contours. Also shown are the point with the minimum velocity RMSE (green), the minimum acceleration RMSE (yellow), and the selected filter design (red).

with better attenuation characteristics. If an alternative way to get a low-lag, linear phase acceleration estimate were available, such as a less noisy accelerometer, then the solution presented here would be different.

3.4.3 Braking Power Calculation

In previous research, methods to calculate the braking power with only a status indication were developed. These have been reviewed in Section 3.2. The previous real-time capable method made use of the difference between a vehicle model's acceleration and the estimated acceleration. As such, errors in the vehicle model led to equal or greater errors in the estimate of braking losses.

When a model's parameters are off, the modeled acceleration is biased, which negatively affects the accuracy of the braking loss estimate ($\hat{P}_{AD}$). This phenomenon is easily seen in Figure 3.20 as the drift of the integrated acceleration from the true vehicle velocity. In the figure, the blue line shows how the modeled velocity, which does not include braking, and

the actual velocity diverge only when the vehicle brakes. The orange line shows that when the model parameters are wrong, an error in the dead-reckoned velocity accumulates. This

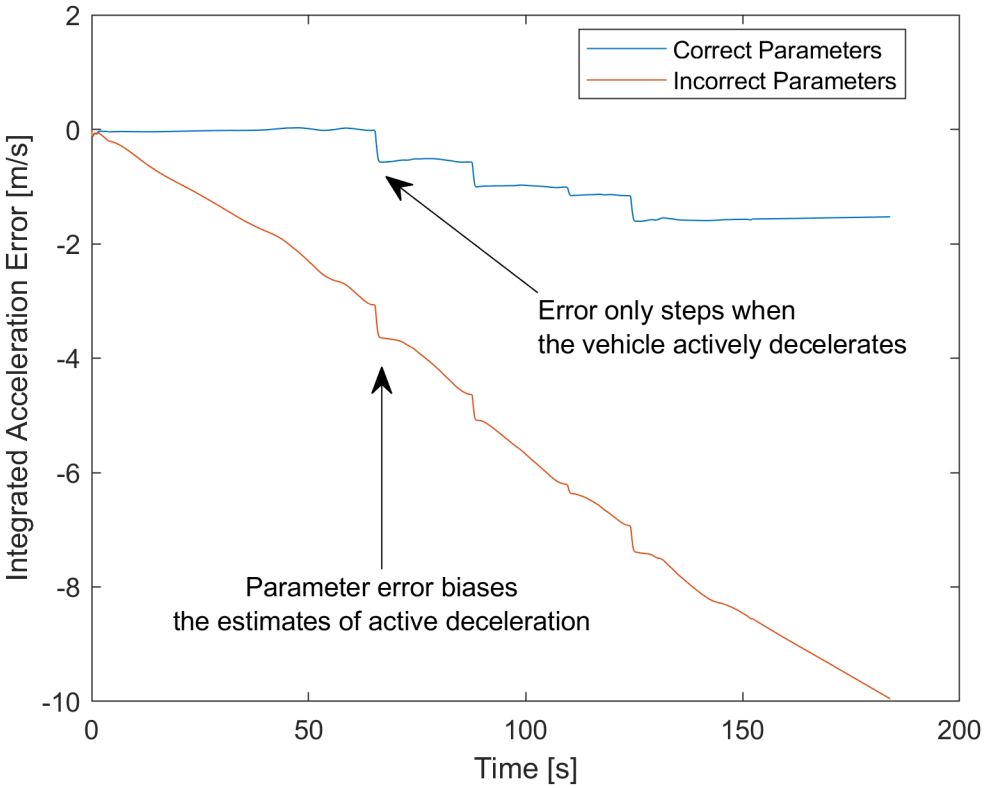

Figure 3.20: Modeled Velocity Drift.

problem is well known when dead-reckoning is used, such as when position is inferred from integrated vehicle kinematics. The model error will bias $\hat{P}_{AD}$ according to the slope of the drift, depending on whether the parameters are too high or too low. Incorrect parameters will have a worse impact when braking lasts longer or is of lower magnitude, as the accumulated error will grow in proportion to the accumulated braking loss.

In theory, the slope of the velocity drift in Figure 3.20 can be learned and used to adjust the braking power estimate. Using the slope to modify vehicle parameters would be an erroneous approach. The impact of an elevated rolling resistance coefficient (C_{rr}) on acceleration is collinear with the impact of an increased drag coefficient ($C_d A_f$) or even an increased effective mass (m_{eff}) in some cases. Therefore, it becomes challenging to discern whether model errors are attributable to one factor or another, as their effects are often

indistinguishable. However, the average error in the modeled and measured acceleration may yet prove useful in increasing the accuracy of the braking loss estimation.

In this work, a new method is introduced, which utilizes the estimated acceleration error to compensate for the impact of inaccurate parameters. This approach closely resembles the one described in [11], but is expected to exhibit greater resilience. While the technique described in [11] relied on point estimates of the mean acceleration error, the novel procedure employs the true mean, thus enhancing its robustness. Two distinct strategies for integrating the mean error are suggested:

- Using a low-pass digital filter of the error, referred to as the Constant Offset routine
- Using Recursive Least Squares (RLS), referred to as the RLS routine

Each of these strategies will be presented next.

The Constant Offset $\bar{P}_{AD}$ Adjustment Routine

First, the Constant Offset routine is given in Algorithm 1. The inputs to the algorithm are the modeled acceleration a_{model}, the true acceleration estimate $\hat{a}$, a boolean flag for braking B. The modeled acceleration is based on a Newton's second law sum of forces, though the braking force is purposely excluded. The true acceleration estimate described in Section 3.4.2 comes from actual vehicle wheelspeed measurements. Finally, the braking boolean B is simply a flag that indicates whether the braking occurred during each time step, regardless of intensity.

Algorithm 1 Adaptive Braking Power Estimation: Constant Offset.

Inputs:
  a_{model} for $1, \ldots, N$
  $\hat{a}$ for $1, \ldots, N$
  $B \in \langle \text{true, false} \rangle$ for $1, \ldots, N$
  m_{eff} for $1, \ldots, N$
  v for $1, \ldots, N$
Parameters:
  $l \triangleq$ Filter window length
  $\Delta_B \triangleq$ Filter update hysteresis
Initialize:
  $\bar{e} \leftarrow 0$
  $i \leftarrow 1$
for $n = 1, \ldots, N$ **do**
  if all$(\neg \left[B(\max(1, n - \Delta_B) \quad \cdots \quad B(n)\right])$ **then** ▷ No braking from n to $n - \Delta_B$
    $\bar{e} \leftarrow \frac{i-1}{i}\bar{e} + \frac{1}{i}(\hat{a}(n) - a_{model}(n))$ ▷ Update $\bar{e}$
    if $i < l$ **then** ▷ Growing window up to l
      i++
    end if
  end if
  $\hat{a}_{brakeless}(n) \leftarrow a_{model}(n) + \bar{e}$ ▷ Always add $\bar{e}$ to a_{model}
  $\hat{a}_{decel}(n) \leftarrow \begin{cases} \hat{a}_{brakeless}(n) - \hat{a}(n) & B \\ 0 & \neg B \end{cases}$
  $\hat{P}_{AD}(n) \leftarrow m_{eff}(n) \cdot \hat{a}_{decel}(n) \cdot v(n)$
end for

The parameters of the algorithm are:

- the window length l, which determines the weight of the current sample in proportion to the past sample average.
- the update hysteresis Δ_B, which controls how many non-braking samples must pass before the updates will resume.

From there, the adaptive task is to compare a_{model} to $\hat{a}$ and estimate the error between the two recursively. The filter output is $\hat{a}_{brakeless}$, which represents an adjustment a_{model}. Ordinarily, Algorithm 1 will simply be updating $\hat{a}_{brakeless}$ recursively through the following equations:

$$e(n) = \hat{a} - a_{model} \tag{3.42}$$

$$\bar{e} = \frac{l-1}{l} e(n-1) + \frac{1}{l} e(n) \tag{3.43}$$

$$\hat{a}_{brakeless} = a_{model} + \bar{e} \tag{3.44}$$

Equation 3.43 reveals that $\bar{e}$ is simply a digitally filtered e. The window length l is allowed to start at 1 and grow indefinitely if memory permits, though $l = \infty$ is not ideal, since the error is bound to spike at times and such samples should be forgotten. In fact, l can be selected to set the cutoff frequency ω_c of the filter, using the relationship in Equation 3.45.

$$\frac{l-1}{l} = \exp\left(-\frac{\omega_c}{f_s}\right) \tag{3.45}$$

The most suitable value of l is influenced by the behavior of $\bar{e}$. For example, for $l = 100$ and $f_s = 10$ Hz, the time constant is $\tau = \frac{1}{\omega_c} = 9.95$ seconds, which would capture the $\bar{e}$ dynamics from the last $4 \times \tau = 40$ seconds. If a particular ω_c or time constant τ is specified, l does not have to be an integer.

Upon braking detection, the Constant Offset algorithm enters a prediction mode, and the average error $\bar{e}$ remains unchanged until braking has not been identified in the previous Δ_B

samples. Incorporating hysteresis through Δ_B helps prevent the unflagged braking dynamics from being transmitted to $\bar{e}$. The calculation of $\hat{a}_{brakeless}$ remains unchanged according to Equation 3.44, with $\bar{e}$ being a constant value. Enhancing the algorithm could involve integrating the estimation of variance and confidence intervals of $\bar{e}$ to prevent inaccuracies in the error estimate from causing significant deviations in $\bar{P}_{AD}$, as well as providing a confidence interval for the entire braking estimate. However, this was not considered in the current scope.

The Recursive Least Squares $\bar{P}_{AD}$ Adjustment Routine

The RLS routine shares many similarities with the Constant Offset routine, but in addition to estimating the constant offset between a_{model} and $\hat{a}$, it also estimates the gain of one versus the other for the previous p samples, where p is the filter order. As such, the RLS routine learns the weights for the constant offset and the previous p samples at times when braking is not occurring. For simplicity, p has been chosen as 1 here. Update equations for the RLS filter have been derived in the literature many times; here, the familiar RLS with a forgetting factor λ is used. See [112] for further details about RLS. The pseudocode for P_{AD} estimation by RLS is given by Algorithm 2.

Demonstration of Adaptive $\bar{P}_{AD}$ Calculation

To demonstrate the way the algorithms work and their limitations, the resulting $\bar{P}_{AD}$ estimate on a simulation is shown in this subsection. A few different cases are examined: first, with correct parameters to demonstrate the clean case, then with incorrect $C_d A_f$, to show how the adaptive algorithm adjusts for drag coefficient errors, then finally with an incorrect vehicle mass parameter. The RLS routine of Algorithm 2 is shown with $\lambda = 0.995$ and order $p = 1$; the Constant Offset routine yields similar results to those shown here.

First, a single simulation was chosen for demonstration purposes. The particular simulation that was chosen had a set speed of 108 km/h (67 mph) and a vehicle mass of 23000 kg.

Algorithm 2 Adaptive Braking Acceleration Estimation: RLS.

Inputs:
a_{model} for $1, \ldots, N$
$\hat{a}$ for $1, \ldots, N$
$B \in \langle \text{true, false} \rangle$ for $1, \ldots, N$
Parameters:
$p \triangleq$ Filter order
$\lambda \triangleq$ Forgetting factor
$\delta \triangleq$ Initial weight for $\mathbf{P}$
$\Delta_B \triangleq$ Filter update hysteresis
Initialize:
$\underset{(p+1)\times 1}{\mathbf{w}} \leftarrow [0, 0 \cdots 0]^\mathsf{T}$
$\underset{(p+1)\times 1}{\mathbf{x}} \leftarrow [1, 0 \cdots 0]^\mathsf{T}$ ▷ Includes an intercept term
$\underset{(p+1)\times(p+1)}{\mathbf{P}(0)} \leftarrow \delta\mathbf{I}$
for $n = 1, \ldots, N$ **do**
$\begin{bmatrix} x(3) & \cdots & x(p+1) \end{bmatrix} \leftarrow \begin{bmatrix} x(2) & \cdots & x(p) \end{bmatrix}$
$x(2) \leftarrow a_{model}(n)$
if all$(\neg \begin{bmatrix} B(\max(1, n - \Delta_B) & \cdots & B(n) \end{bmatrix})$ **then**
$\alpha \leftarrow \hat{a}(n) - \mathbf{x}^\mathsf{T}\mathbf{w}$
$g \leftarrow \mathbf{P}\mathbf{x}(\lambda + \mathbf{x}^\mathsf{T}\mathbf{P}\mathbf{x})^{-1}$
$\mathbf{P} \leftarrow \lambda^{-1}\mathbf{P} - gx^\mathsf{T}\lambda^{-1}\mathbf{P}$
$\mathbf{w} \leftarrow \mathbf{w} + \alpha g$
end if
$\hat{a}_{brakeless}(n) \leftarrow \mathbf{x}^\mathsf{T}\mathbf{w}$
$\hat{a}_{decel}(n) \leftarrow \begin{cases} \hat{a}_{brakeless}(n) - \hat{a}(n) & B \\ 0 & \neg B \end{cases}$
$\hat{P}_{AD}(n) \leftarrow m_{eff}(n) \cdot \hat{a}_{decel}(n) \cdot v(n)$
end for

White noise with a variance of 0.02 m/s and 0.5 degrees was introduced into the simulation's wheelspeed and grade, respectively. Noise addition was performed before the calculation of $\hat{a}$ and a_{model}.

For the first case, the correct parameters were used to calculate a_{model}. Figure 3.21 outlines the calculation of $\bar{P}_{AD}$ for this correctly parameterized case. Vertical dashed lines denote the onset and offset of braking events. Starting with the top panel, it can be seen that the measured acceleration (in yellow) tracks the true acceleration (in black), even during braking events. In contrast, the modeled acceleration does not track the true acceleration

during the braking events, which is by design. The unadjusted a_{model} (blue line) and the adjusted a_{model} (orange line) display this separation from true acceleration during braking events. Moving to the middle panel of Figure 3.21, the error in braking deceleration (a_{decel}) is displayed, which is calculated as $\hat{a}_{decel} - a_{decel,\,true}$. The true braking deceleration $a_{decel,\,true}$ is derived from the true braking force, which is available because this is a simulation. Ideally, the traces of the middle panel would have zero mean and minimal variance, especially during braking events.

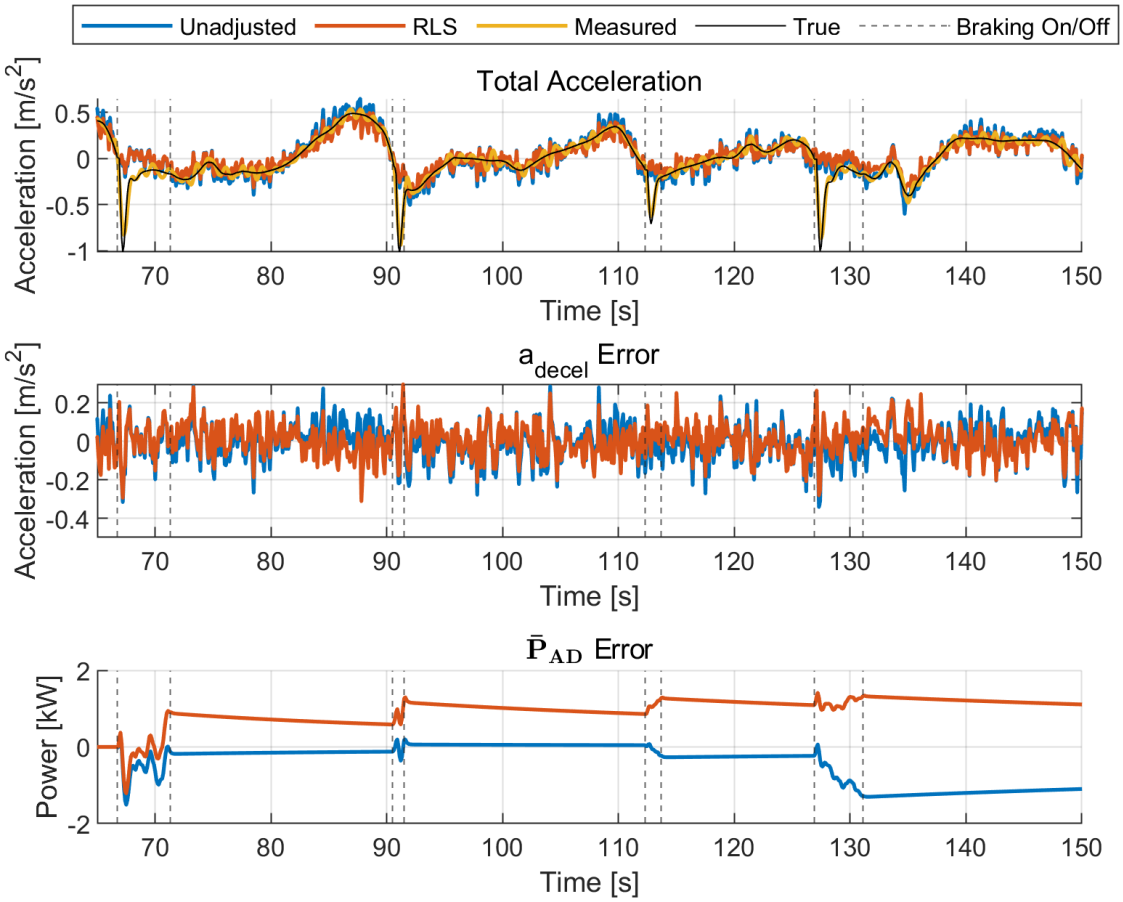

Figure 3.21: Results of Algorithm 2 with correct parameters. The onset/offset of braking events are marked by gray vertical lines.

Finally, the lower panel of Figure 3.21 shows the resulting error in average braking power ($\bar{P}_{AD}$), which is the true metric of interest. It can be seen that during each braking event, the error in the braking power shifts. In between braking events, the error decreases as the rolling average dilutes the error. There is no clear sign that the RLS adjustment performs better or

worse than the unadjusted braking power. The absence of a clear advantage is anticipated because the unadjusted braking applies the appropriate parameters in this instance and, therefore, does not require any modifications.

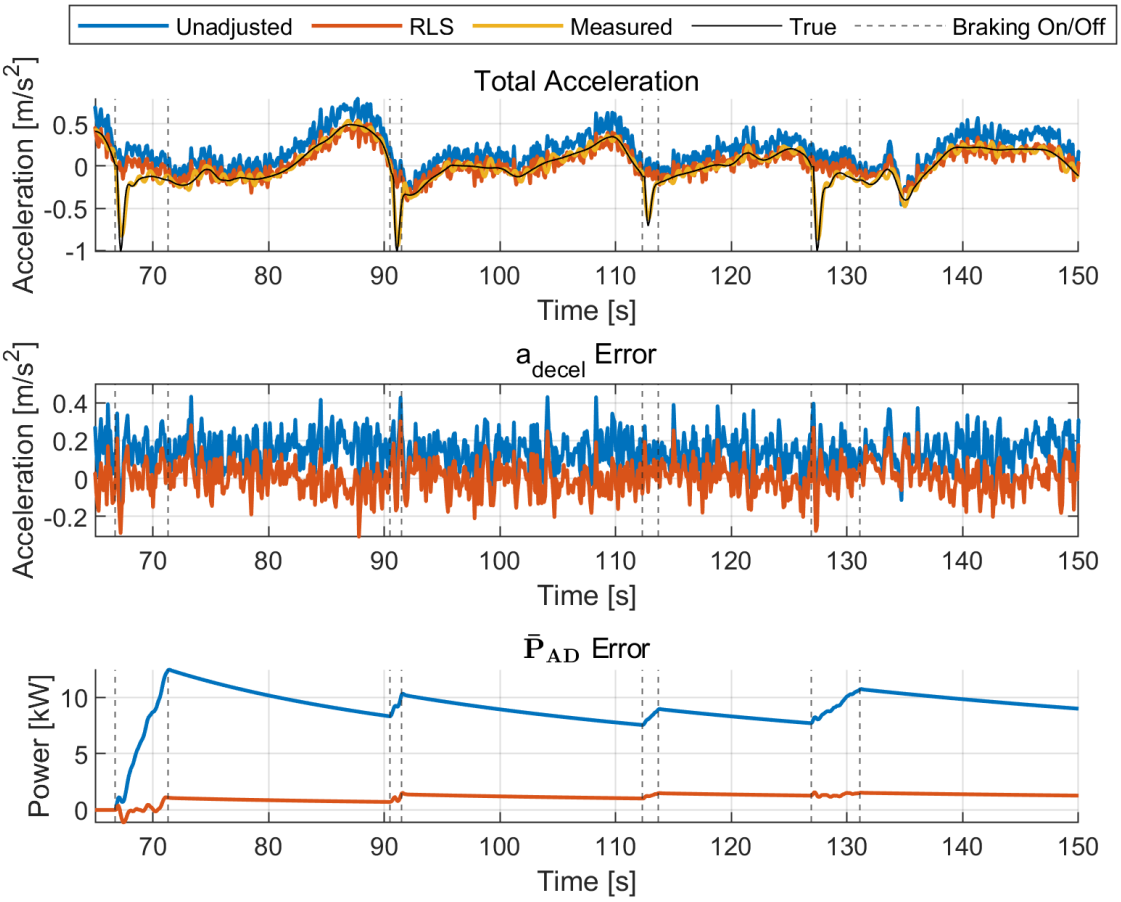

Figure 3.22: Results of Algorithm 2 with incorrect drag and rolling resistance coefficients (Error in $C_dA_f = -5$ m^2, error in $C_{rr} = -0.005$). The onset/offset of braking events are marked by gray vertical lines.

Next, in Figure 3.22, the drag and rolling resistance coefficients were incorrectly specified. Both were shifted drastically below the actual value to better emphasize differences (C_dA_f was shifted by -5, and C_{rr} by -0.005). Because both coefficients being made lower than the true value, the modeled acceleration is consequently shifted above the actual observed acceleration, which RLS works to correct. In the middle panels, which show the error in a_{decel}, it can be seen that the RLS errors (orange) are closer to zero than the unadjusted errors (blue). This leads to a better estimate in $\hat{P}_{AD}$, as shown by the lower panel of Figure

3.22; the error in $\bar{P}_{AD}$ by the end of the short simulation is 1.28 kW with RLS adjustment versus 9.04 kW without adjustment. These represent errors of 8% versus 56% in the true average braking power.

Finally, we look at how the $\bar{P}_{AD}$ estimates are affected by incorrect mass. Mass was incorrectly specified as 28000 kg instead of 23000 kg, a 22% error. Figure 3.23 shows that if the mass is incorrectly specified, a_{decel} can be well tracked, but the estimate for P_{AD} will be wrong. In fact, it appears that the RLS-adjusted $\bar{P}_{AD}$ estimate is more sensitive to mass errors than the unadjusted estimate of $\bar{P}_{AD}$, based on the increased error of the orange trace in the lower panel of Figure 3.23.

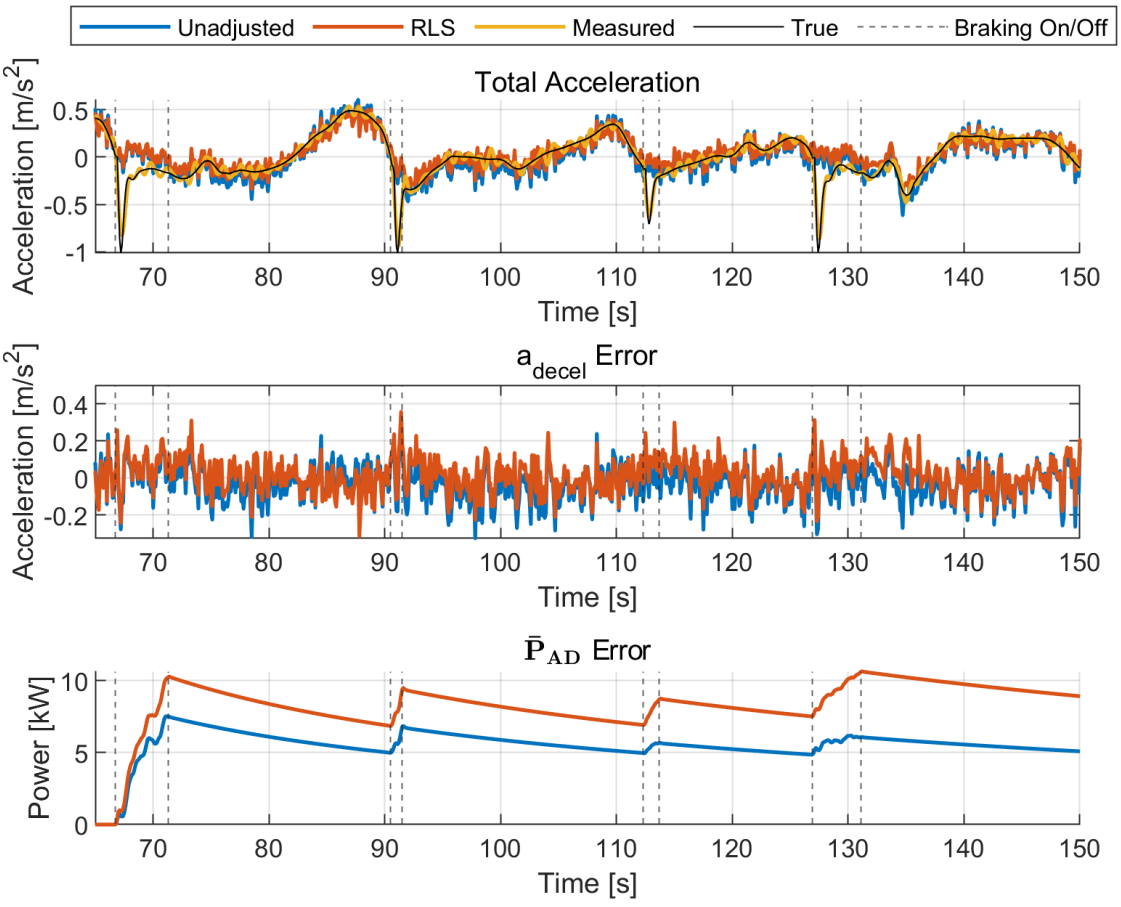

Figure 3.23: Results of Algorithm 2 with incorrect mass (Error in $m = +10000$ kg). The onset/offset of braking events are marked by gray vertical lines.

It has been demonstrated that the impact of measurement noise and model errors on braking power estimates can be ameliorated by using the adaptive RLS algorithm, even with

large errors in the drag and rolling resistance coefficients. Mass errors are not taken into account by the adaptive algorithm, so including adaptive vehicle mass estimation has been left to future work.

This subsection represents a first view of the characteristics of the braking power estimate. It is important to emphasize the broad applicability of the method. The only measurements required were vehicle mass, velocity, an acceleration model, and a Boolean indicator for braking. As such, it is highly portable to a wide variety of vehicle architectures. The sensitivity and robustness of these two novel adaptive brake power estimation routines will be reviewed in the results section.

This concludes the section focused on the practical implementation of the platoon power estimation framework. While aerodynamic drag reduction modeling was discussed, additional focus was given to the estimation of braking power ($\bar{P}_{AD}$), as this method does not exist in the literature. All proposed elements are real-time capable, with the slight caveat that the proposed linear phase acceleration filter runs at a small delay. This delay could be easily reduced. In the next section, which concludes Chapter 3, the top-level validation methodology will be described, including a description of the experimental data and the data analysis methodology.

3.5 Validation Methodology

Now that practical methods for implementing the platooning energy consumption quantification have been proposed, we turn attention to another question: How good is the method at quantifying the energy consumption of real platoons? If the framework proposed in this dissertation is deemed to be insufficiently precise, it is crucial to pinpoint the reasons for the lack of precision. Earlier in the chapter, it was shown that $NPC_{inferred}$ and $NFC_{inferred}$ accurately reflect the variations in platoon energy consumption when the vehicle parameters

and actual road loads are known. Validation is essential to identify instances where inaccurate vehicle parameters can result in incorrect assessment of platoon energy usage. To perform the validation, two approaches will be used.

1. Using the simulation results, do a sensitivity analysis on errors in modeled truck parameters, $\hat{P}_{AD}$, and $PR\hat{R}_{DR}$.
2. Using experimental platooning data from different campaigns with different speeds, masses, trucks: apply the methodology, and as far as possible, see how well it describes the measured differences in platoon energy consumption (e.g. NPC_{true}).

The types of errors to be investigated using the simulation results are:

- Measurement noise:
  - Velocity
  - Grade
  - Engine Power
- Parameter errors:
  - Coefficient of rolling resistance C_{rr}
  - Modeled truck mass m
  - Drag area C_dA

For the experimental data, there are many sources available, each with its nuances. Some are test track results, others are real applications, some over hilly terrain, others on flat ground. Some of the datasets meet all the model assumptions, while others do not, involving changes in elevation, unknown parameters, etc. As a result, the validation power of each experiment is unique. In the following subsection, each experiment will be described in detail.

3.5.1 Experimental Datasets

Four experimental data sets were available for platooning. All of the experimental data uses the same four trucks, with varying amounts of trailer load. The truck specifications that are common between the datasets are given in Table 3.6. There were two different styles of trucks, the MY2015 sleeper cabs A1 and A2, and the MY2009 military daycabs T13 and T14. All four trucks featured automatic/automated transmissions and inline six engines in the 12-15 liter displacement range.

Table 3.6: Specifications of the trucks used in the Datasets.

Truck ID	A1	A2	T13	T14
Model Year	2015	2015	2009	2009
Manufacturer	Peterbilt	Peterbilt	Freightliner	Freightliner
Model	579	579	M915A5	M915A5
Engine	Paccar MX-13	Cummins ISX15	DDEC S60	DDEC S60
Transmission No. Speeds	Eaton Fuller 10	Eaton Fuller 10	Allison 4500 6	Allison 4500 6
Steer Tires Cold Pressure	Michelin XZA3 120 psi	Michelin XZA3 120 psi	Michelin XZUS 130 psi	Michelin XZUS 130 psi
Drive Tires Cold Pressure	Michelin XDA 90 psi	Michelin XDA 90 psi	Michelin XZUS 90 psi	Michelin XZUS 90 psi
Trailer Tires Cold Pressure	Various[1] 100 psi	Various[1] 100 psi	Various[1] 100 psi	Various[1] 100 psi
Peak Torque at	2372 Nm (1750 ft.lbs) 1000 RPM	2237 Nm (1650 ft.lbs) 1000 RPM	2237 Nm (1650 ft.lbs) 1200 RPM	2237 Nm (1650 ft.lbs) 1200 RPM
Claimed Power	320 kW (430 hp)	309 kW (415 hp)	373 kW (500 hp)	373 kW (500 hp)
C_dA_f	5.5[2]	5.5[2]	7.0	7.0
C_{rr}	0.0055	0.0055	0.0055	0.0055

[1] Trailers varied throughout tests and occasionally had a multiple models of tire mounted

[2] Except during NRC testing (Dataset 2), where 5.0 is a better estimate due to the boat tails on the trailers

The four datasets are:

1. Test-track data from the DOE-funded project titled "Fuel-Efficient Platooning in Mixed Traffic Environments", which was collected from 2019 to 2021 at the NCAT test track in Opelika, AL and the ACM test track in Ypsilanti, MI. These data include truck platoons of up to four members. This dataset includes asynchronous baselines.
2. Test-track data from a 2019 platooning campaign on the PMG test track in Blainville, Canada, which formed the basis of [31, 42, 34, 33] among others. The trials are 1 hour long, and the trucks were very well instrumented. This dataset includes synchronous baselines.
3. On-road data from October 2019 two-truck platooning on Interstate 85 which was analyzed in [73]. Baselines are asynchronous and trials are about 40 minutes long.
4. On-road data from April 2021 platooning trials on AL Hwy-280 that was analyzed in the journal publication [48]. This dataset includes synchronous baselines, but each trial is only 20 minutes long.

Next, each of the four datasets will be described in more detail. This is intended to give a sense of the breadth and depth of the data contained in the four datasets, since ultimately the data will be used to validate the proposed energy prediction framework.

Dataset 1: DOE Test Track

Dataset 1 was collected at the National Center for Asphalt Technology (NCAT) Test Track in Opelika, AL and the American Center for Mobility Highway Loop in Ypsilanti, MI between the fall of 2019 and the spring of 2021. The motivation of the study was to gain a deeper understanding of the performance of platoons in the presence of disturbances and to improve this performance.

The testing was carried out in two phases. In the first phase, the baseline energy consumption of the two- and four-truck platoons was established on the NCAT track, which is flat, and the ACM track, which is hilly. In this case, the baseline was a classical PID platoon

IVD controller that was not designed to traverse hills or handle disturbances efficiently. In the second phase, the trucks were tested again on the flat track and the hilly track with an improved gap controller[13]. Although the study did not follow the SAE J1321 standard procedure, the spirit of the J1321 protocol was observed to control unwanted variation, including the following:

- Trucks were warmed up 1 hour before testing could commence, and downtime between tests was limited to 30 minutes
- Each trial lasted 1 hour, for a total of 26 laps at NCAT and 19 laps at ACM.
- Regeneration of diesel particle filter (DPF)[14] was manually forced outside of the testing.
- Rain and persistent winds were grounds for invalidation of the test.
- Climate control and accessory loads were fixed for each test.

Four trucks were used for the study, including two 2015 Peterbilt 579 sleeper cabs (A1 and A2) and two Freightliner M915 day cabs (T13 and T14), one of which is armored for military duty (T13). Table 3.7 lists the test details for Dataset 1. The track limitations restricted the test speed to 20 m/s (45 mph), which is lower than the usual speeds on highways, leading to a reduced drag fraction. In order to enhance the significance of the outcomes, the trailers were left unloaded, causing an increase in the drag fraction for the trucks. As a result, the drag fraction during the tests was around 60%, making it relevant for highway conditions.

Two different two-truck platoons and a four-truck platoon were tested. For the two-truck platoons, A1 led T14 , or T13 led A2 . The four-truck platoon configuration consisted of these concatenated two-truck platoons, starting with A1 , followed by T14 , T13 , and finally A2 .

Two different gap control strategies were utilized in Dataset 1. For phase 1 testing, the objective was to collect baseline performance data for platoons with increasing levels of

[13]This improved gap cotroller was validated by Dataset 4 and its accompanying publication [48]

[14]DPF regeneration burns excess soot in the exhaust after-treatment system. The process increases fuel consumption and usually occurs invisibly to the driver.

Table 3.7: Specifications of Dataset 1 in addition to those in Table 3.6.

	A1	A2	T13	T14
Truck+Trailer Weight	16,180 kg (35,660 lbs)	17,250 kg (38,020 lbs)	22,890 kg (50,460 lbs)	19,510 kg (43,020 lbs)
Testing Dates	Sept.-Nov. 2019; Jul.-Sept. 2020; May 2021			
Testing Location	NCAT and ACM test tracks			
Test Speed	20.1 m/s (45 mph)			
Test Duration	1 hour each trial			
Platoon Configurations	A1⇐T14; T13⇐A2; A1⇐T14⇐T13⇐A2			
Gap Control Strategies	PID (phase 1) and Lookahead (phase 2)			
Gap Setpoints	10.6, 15.2, 22.9, and 30.5 meters (35, 50, 75, and 100 feet)			
Baselines	Yes, Asynchronous			
Trailer Aero	Box trailer with skirts[1]			
Baselines	Yes, Asynchronous			
Hills	Yes, at ACM			
Cut-ins	Yes, between T14 and T13 in the four-truck platoon			
Surrounding Traffic	Only that due to cut-ins			

[1] All trailers were skirted except T14 ACM phase 2, where no skirted trailer was available for rental.

external factors. As noted in by literature review, the platoon gap controller is a crucial factor in how a platoon deals with external factors. Therefore, the baseline in this case was the PID-based IVD control strategy of [54]. In phase 2 of Dataset 1, PID control was traded for lookahead control, allowing the platoon to traverse hills much more efficiently.

The IVD setpoints tested for each platoon also varied, from a minimum of 35 feet (10.6 m) to a maximum of 100 feet (30.5 m). It should be noted that while the PID gap controller used in Phase 1 kept the gap setpoints rigidly, the Phase 2 lookahead controller allows flexible shrinkage and growth of the gap by design, in its effort to balance gap control with fuel efficiency. Therefore, the IVD setpoints were nominal for the lookahead controller, although on average the lookahead control strategy resulted in an IVD within 10 feet of the set value.

The reference performance of the trucks in Dataset 1 consists of each truck driving alone with cruise control. The cruise controller is capable of using the engine brake, and sometimes did, so the reference runs are not strictly brakeless. Due to testing constraints, reference runs had to be collected asynchronously, introducing additional uncontrolled variation.

The trailers in Dataset 1 were all standard box trailers with skirts. During the ACM 2021 testing, T14's rental box trailer arrived without skirts, but it is difficult to ascertain the effect this had on the results.

Regarding hills, only the ACM test track has significant hills that could pose a challenge to platoons. The hills at ACM are shorter and steeper than the national average [59, 10], leading to more rapid and intense disturbance dynamics on the platoon gap controller. The NCAT track only has slight elevation changes, with one straight going slightly uphill and the other straight going slightly downhill.

Cut-in tests were conducted during both phases. The only platoon configuration subjected to cut-ins was the four-truck configuration. The cut-in vehicle was inserted between the second and third trucks, T14 and T13. In some cases, the vehicle remained until a straight was reached, and, in others, the vehicle cut out briefly after cutting in. The particular behavior of the cut-in vehicle was subject to the driver, which varied between phases. The author conducted an analysis of the cut-ins in [69], finding that the cut-ins did increase the fuel consumed by the third and fourth trucks appreciably.

Due to the cyclical nature of the test track trials, assumptions about the braking energy manifested as power consumption (Assumption 3) are satisfied. With asynchronous baselines, it is feasible to calculate the NFC_{true} of the Dataset 1, although the variation in reference fuel consumption will introduce an error in the target variable, violating the assumptions of ordinary least squares.

The truck masses for Dataset 1 were weighed using CAT scales, but are subject to small variations as the amount of fuel, the number of occupants and the equipment in each truck changed. Furthermore, the trucks did not have the same trailers for all four trials (NCAT

and ACM in phases 1 and 2): Only two trailers were owned, so at least two trailers were rented for each trial. The road load parameters of the truck are approximate, especially for T13 and T14. An aerodynamic and rolling resistance analysis was conducted on Peterbilt trucks (A1 and A2) in [33], which gave good certainty about C_dA and C_{rr}.

Overall, the ability of Dataset 1 to validate the model is very good compared to the other datasets. The sheer amount of data is high, and the data is reasonably well controlled, with the main drawback being the lack of a synchronous baseline. One distinguishing feature of the dataset is that it includes large variations in both braking and drag reduction. As a result, Dataset 1 will be one of the most powerful validation sets.

Dataset 2: NRC/Auburn/NREL at the PMG Test Track

The second of the datasets here has already been partially introduced in the cut-in case study (Section 3.2). It will be interesting to see whether $NPC_{inferred}$ generalizes better than the previous methodologies that were applied in Section 3.2. Dataset 2 is well described in a series of publications [31, 42, 33], so what follows is a brief summary.

For Dataset 2, only trucks A1 and A2 were used, together with a control truck. The 65 mph, 65000 lb load condition of the trucks is highway relevant. With gravimetric fuel results, dedicated on-site test support personnel, weather station measurements, and a control truck, the test comes closest to being a true SAE J1321 test out of any of the datasets in this dissertation.

IVD gap control was performed via PID control, with the NMPC controller being unnecessary for the flat-track setting. Dataset 2 was performed with a finer resolution of IVDs than the others but contains only one braking test, the cut-in case. Therefore, Dataset 2 will be the most useful to assess the suitability of the estimates of $\bar{P}_{aero}$. Six IVD setpoints were tested, ranging from quite close at 9 meters to a distance of 3 seconds, 79 meters, representative of a stock ACC IVD.

Table 3.8: Specifications of Dataset 2 in addition to those in Table 3.6.

	A1	A2	T14	T13
Truck+Trailer Weight	29,500 kg (65,000 lbs)	29,500 kg (65,000 lbs)	- -	- -
Testing Dates	June-July 2019			
Testing Location	Transport Canada Motor Vehicle Test Center "Bravo" Track			
Test Speed	29 m/s (65 mph)			
Test Duration	13 laps (52 miles), from key-on to key-off			
Platoon Configurations	A1⇐A2			
Gap Control Strategies	PID			
Gap Setpoints	9.1, 12.2, 15.2 22.9, 45.7, 78.6 meters (30, 40, 50, 75, 150, 258 feet)			
Trailer Aero	Box trailer with skirts and boattails			
Baselines	Yes, Synchronous			
Hills	No			
Cut-ins	Yes, 1 trial			
Surrounding Traffic	Yes, several trials			

Another distinguishing feature of Dataset 2 is that the trailers were equipped with boat tails, increasing the benefits of platooning. The drag reduction ratio is likely to be higher for the Dataset 2 than for the others. It should be noted that rear trailer aerodynamic devices are not widely adopted compared to skirts as of 2021 [113].

Dataset 3: Interstate 85 Long Trials

The first of the on-road datasets discussed here comes from on-road trials on Interstate 85. This dataset was collected during the same time and in a very similar spirit to Dataset 1, although it only used Peterbilt sleeper-cab tractors (A1 and A2). The trial was carried out on Interstate 85 from Exit 60 south to Exit 38 and back (see Figure 3.24). The data sometimes includes the drive from the NCAT track to I85, which does not belong in the final analysis, so any data with latitude and longitude exceeding 32.6435 and -85.3518 was removed.

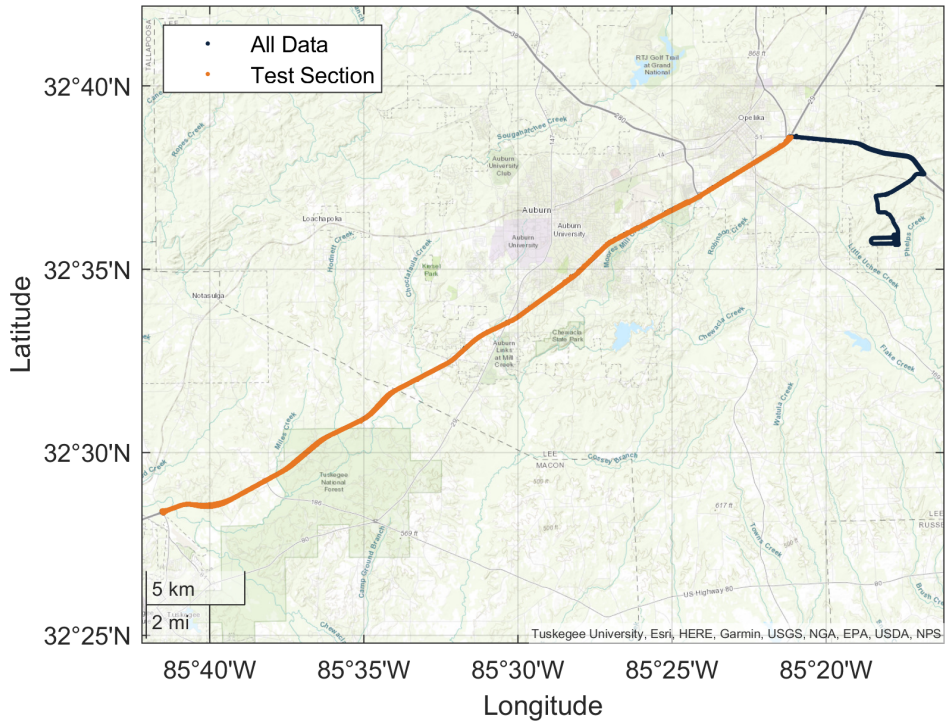

Figure 3.24: Map of the GPS data of Dataset 3.

Both A1 and A2 were as they were during Dataset 1, since both campaigns were conducted on the same days. Since Dataset 3 was conducted on-road, the test occurred at 65 mph. Like Dataset 1, the baselines were asynchronous. The sole platoon consisted of the two Peterbilts, A1 followed by A2, using the PID gap control strategy. The trials are approximately 45 minutes in total length, with the loop being only 45 miles long.

Two different IVDs were tested: 75 feet and 150 feet, representing time gaps of 0.78s and 1.57s, respectively. Both IVDs are closer than the 2-second following distance that HDV drivers naturally follow at [74], and may be propoerly considered "platooning". There were a total of seven 75-foot runs, five 150-foot runs and 4 and 5 baseline reference runs for trucks A1 and A2 respectively.

Due to its out-and-back nature, Dataset 3 contains two separate platooning microtrips per trial, first the microtrip from Exit 60 to Exit 38, then back. On the first microtrip, the elevation decreases from 220 to 80 m, and on the second microtrip, the trend reverses.

Table 3.9: Specifications of Dataset 3 in addition to those in Table 3.6.

	A1	A2	T14	T13
Truck+Trailer Weight	16,180 kg (35,660 lbs)	17,250 kg (38,020 lbs)	- -	- -
Testing Dates	Sept.-Oct. 2019			
Testing Location	Interstate 85 between Alabama Exits 60 and 38			
Test Speed	29 m/s (65 mph)			
Test Duration	45 minutes each trial			
Platoon Configurations	A1⇐A2			
Gap Control Strategies	PID only			
Gap Setpoints	22.9 and 45.7 meters (75 and 150 feet)			
Trailer Aero	Box trailer with skirts			
Baselines	Yes, Asynchronous			
Hills	Yes			
Cut-ins	Yes, from real traffic			
Surrounding Traffic	Yes, unquantified			

The explanatory power of Dataset 3 is hampered by real-world variation. The lack of a synchronous baseline and many other variations will make it difficult to accurately estimate the NFC. Even if there were many more replications, there is still no precise knowledge of surrounding conditions and the effect thereof. Therefore, it remains to be seen how much Dataset 3 can be used to validate the proposed platoon energy consumption model, but it may also represent the most realistic look at practical platooning performance.

Dataset 4: AL Highway 280 Lookahead Trials

The last of the four datasets was collected on AL Highway 280 over two days in April of 2021. This dataset was previously analyzed in [48]. The experiment was structured as an abbreviated SAE J1321 test, where each "run" was only 20 minutes, an out-and-back stretch on US HWY-280. A platoon of T13 and T14 was sent out, followed by a control truck at an IVD greater than 500 meters.

Table 3.10: Specifications of Dataset 4 in addition to those in Table 3.6.

	A1	A2	T14	T13
Truck+Trailer Weight	-	-	22,890 kg	19,510 kg
	-	-	(50,460 lbs)	(43,020 lbs)
Testing Dates	Apr. 12-13, 2021			
Testing Location	Alabama Hwy 280			
Test Speed	24.6 m/s (55 mph)			
Test Duration	20 minutes each trial, excluding turnaround			
Platoon Configurations	T13⇐T14			
Gap Control Strategies	PID, String-stable PID, Lookahead			
Gap Setpoints	35 meters (115 feet)			
Trailer Aero	Box trailer with skirts			
Baselines	Yes, Synchronous			
Hills	Yes			
Cut-ins	No			
Surrounding Traffic	Yes, unquantified			

The highway loop featured much shorter and steeper hills on the outbound (eastward) section and longer rolling hills on the way back to the starting point (inbound). The outbound section also featured a net negative elevation change. As such, the control strategy was much more taxed on the outbound section. This was seen as an excellent opportunity to compare controllers in different grade circumstances.

Further details may be found in the source publication, but, in general, classical feedback control was compared to a novel decentralized predictive controller. Two types of classical control were performed, but the performance difference between the classical controllers was trivial compared to the benefits of predictive control. It was discovered that the classical control could not achieve savings on the outbound section, but that savings were identical on the inbound section.

All tests were carried out at 55 mph (24.6 m/s) with the IVD set at 35 meters (115 ft, 1.4 s). Standard unloaded trailers with side skirts were used. There were no interruptions

in the platoon during the tests[15]. Unfortunately, the testing constraints resulted in only one replicate of single-truck driving, so there is not much data to ascertain the non-platooning performance of the trucks.

Dataset Summary

At this point, all four of the datasets have been described. Recalling how the literature review explored how internal factors and disturbances affected platooning savings, there is a good breadth of both across the four datasets. A variety of IVDs, hill intensity, traffic interactions, and platoon configurations are encapsulated between the four. Having described what is contained in the datasets, the data analysis procedure for the data will be generally described next.

3.5.2 Data Processing

The data utilized in this dissertation was primarily collected using the Auburn University CACC architecture, which utilizes Robotic Operating Software (ROS) to collect and control the platoon. The platoon data is logged using ROS's `rosbag` feature, which is simply a way of recording the various signals that are being used to operate the "robots", in this case platooning trucks. The system architecture is shown in Figure 3.25. A central computer is responsible for uniting the CAN bus, V2V, Radar, and GPS nodes, and sending the corresponding control commands.

The data is parsed into a non-time-aligned MATLAB structure, which is interpolated linearly for continuous signals or by nearest for discrete signals, such as "On/Off" or current gear. The interpolation distinction is made to avoid issues such as non-sensical gear numbers (there is no 4.97th gear).

For each dataset, all data was concatenated into one table, with each row representing an observation at 10 Hz. Metadata such as the truck, the number of trucks, the IVD set

[15]Possibly because the operation was slower than the speed limit or because there was a fear of cutting between two mil-spec vehicles.

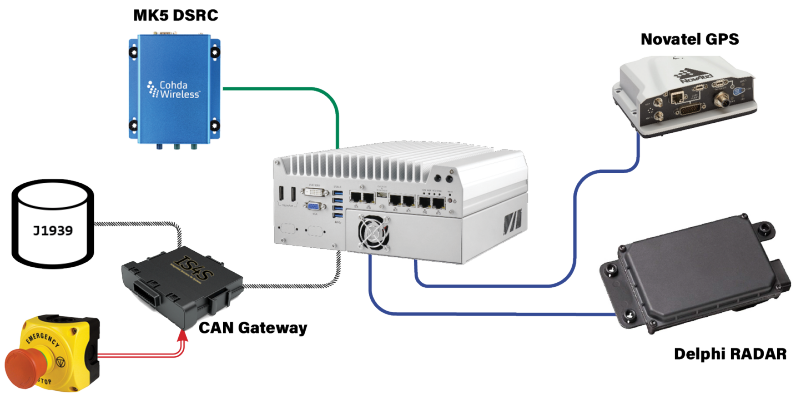

Figure 3.25: Auburn CACC System Architecture.

point, and the iteration number were extracted from the organized filename strings (ex: "A1_2T_100_1_xxxx" is truck A1, in a two-truck platoon, at an IVD of 100 feet). The year of the test and whether the truck was a leader were also added. Although this results in a much larger table "memory-wise" than necessary, it made the analysis easier and less prone to identification errors.

All tables from each run were vertically concatenated into one large table after the interpolation step. Then the `findgroups` function could be used to find matching runs based on metadata. With the availability of a GPS signal, the runs were synchronized with each other using GPS seconds. This is a necessary step to add the IVD range estimate to the leaders, which is required for the *DRR* models. At this stage, Datasets 3 and 4 were segmented into the outbound and inbound portions. It is also useful to transform latitude and longitude into the East-North-Up (ENU) frame.

The Auburn CACC system includes a grade estimator, which uses GPS velocity components to estimate road grade, but it is not a perfectly reliable output, and some of the trucks completely lacked the grade estimate. Because each dataset had been formed into one large tidy table, constructing a grade lookup table was easy.

To acquire a more robust grade lookup, an entire dataset's table was sorted by either integrated wheelspeed or the robust lap position metric developed in [29, 10]. Missing rows were removed, and a moving median of the grade estimate taken. To avoid oversmoothing,

a relatively small window was used for the moving median, depending on the dataset. The lookup table was then used to add the road grade to all trucks, even those with no grade estimate originally. Grade lookup utilities do exist for applications such as this, such as Google's Road Elevation and Snap to Road APIs. These were experimented with, but it was found that the Google grade lookup disagreed rather substantially with the estimated lookup (and gave erroneous results), so it was not used.

Finally, the required features were added to the table, such as $\bar{P}_{AD}$ and $\bar{P}_{aero}$. After the addition of features, the tables were ready for analysis. The features added included:

- Estimated acceleration using the FIR filter $\hat{a}$
- Estimated terrestrial winds and environmental conditions from OpenWeatherMap
- Various drag reduction ratios of [28, 93]: Schmid's [28], Hussein's Rational Polynomial [93], and Hussein's Power Law [93]
- Modeled acceleration a_{model} in three forms: unadjusted, adjusted by RLS, and adjusted by Constant Offset
- Estimated braking deceleration and power, $\hat{a}_{decel}$ and $\hat{P}_{AD}$, also in three forms
- Estimated body-axis wind velocity and yaw, using the OpenWeatherMap API
- Estimated fan power consumption

Care was taken to ensure that the added features were added in the same way that they would be in real time, as opposed to postprocessing. For example, the MATLAB function `filter` was used instead of the function `filtfilt`, as the latter is non-causal and cannot be performed online. The one exception to this rule was the acceleration estimate, which was backshifted to account for the known lag. This means that the filter was running "in the past", albeit by a slight 1.4 seconds.

3.6 Chapter Summary

The chapter began with an overview of the fundamentals of platoon energy consumption, leading to the formulation of assumptions and a discussion on normalization of energy

consumption. This discussion also touched on what the energy consumption should be normalized against, i.e. the choice of reference. Subsequently, an evaluation of existing methods for comprehensive modeling of platoon energy consumption was conducted, drawing on the author's earlier work. Existing methods did not exhibit the desired level of generalizability, prompting a more thorough investigation of the mechanisms that influence platoon fuel consumption. A simulation study was carried out, revealing that normalized power consumption could be deduced from a vehicle's braking losses and drag reduction without the necessity of a reference. The focus then turned to practical strategies for computing each of the power inference subcomponents, such as aerodynamic drag reduction, vehicle acceleration, and braking loss. A technique was suggested to address the impact of parameter inaccuracies on braking loss. Finally, the general approach to validate the framework was described. The chapter culminated in a comprehensive explanation of the four datasets to be utilized and the techniques for processing the data.

Chapter 4

Results and Discussion

Let us summarize the progress made in this dissertation so far. Platooning represents a specific form of vehicle autonomy in which vehicles travel closely together, leading to savings in aerodynamic drag. However, the attractiveness of platooning has diminished in recent years due to the failure to realize the anticipated benefits in practical trials, often due to braking inefficiencies. Although various models have been suggested to reduce drag in platoons, there is still no comprehensive model that combines drag reduction and controller-induced braking losses. Such a model could potentially enhance platooning by providing operators with feedback on their energy savings and the extent of those savings. The proposed methods enable platoon operators to deduce these advantages independently, without external references. This method of inference is formalized as inferred normalized power or fuel consumption, $NPC_{inferred}$ and $NFC_{inferred}$, respectively.

This chapter examines the accuracy and precision of inferred platoon energy savings in two ways. First, a sensitivity analysis of the IPG TruckMaker simulations introduced in Section 3.3.1 will be performed. Second, the experimental results will be analyzed using a pairwise comparison. The sensitivity analysis will aid the comprehension of the experimental results, where a true measurement of the quantities involved was not possible.

4.1 Sensitivity Analysis

We aim to assess the sensitivity of $NPC_{inferred}$ to signal noise and parameter errors. Simulation offers the advantage of having access to actual values, allowing a straightforward assessment of the sensitivity to noise and errors. In our case, the noise and errors in question are those that impact the inferred energy benefits, in the form of $NPC_{inferred}$ or $NFC_{inferred}$ (Equations 3.21b and 3.24). Recall that total power $\bar{P}_{total}$, braking power $\bar{P}_{AD}$, aerodynamic drag power $\bar{P}_{aero}$, and drag reduction ratio DRR are the basic components of $NPC_{inferred}$ and that each of these terms is vulnerable to various errors in noise and parameters.

4.1.1 Noise Characteristics

First, we investigate the appropriate level of noise to apply to wheelspeed, grade, and engine power. Based on the experimental data, wheelspeed encoder noise is on the order of $\sigma_v = 0.01$ m/s, although it is heavier-tailed than in a normal distribution[1]. The empirical wheelspeed error distribution was constructed by subtracting the smoothed wheelspeed from the raw wheelspeed signal. For an explanation of the construction of empirical distributions, see [115]. The wheelspeed encoders also have a small Constant Offset from the true longitudinal velocity, but this is neglected here. GPS velocity can readily correct for this bias in practice.

Also based on the experimental data, the noise in the grade estimate is on the order of $\sigma_\theta = 0.1^\circ$, and is also heavier-tailed than a normal distribution. Furthermore, the GPS grade estimate is subject to both signal dropouts and rapid variance inflation as the view of the various satellites changes. These signal dropouts and error spikes may have a worse effect on the estimate than simple zero-mean noise, but here we will assume that a noisy but reliable grade estimate is available.

The empirical probability density functions of wheelspeed and grade noise are illustrated in Figure 4.1. Also shown in Figure 4.1 are the best-fit T-distributions using maximum likelihood estimation. The shape parameter ν of the T-distribution in both cases is approximately equal to two, whereas the shape parameter of a normal distribution is $\nu = \infty$. This means that wheelspeed and grade error will be more prone to frequent outliers than a normal distribution of the same variance.

[1]For explanations of the normal, uniform, and T-distribution, see [114].

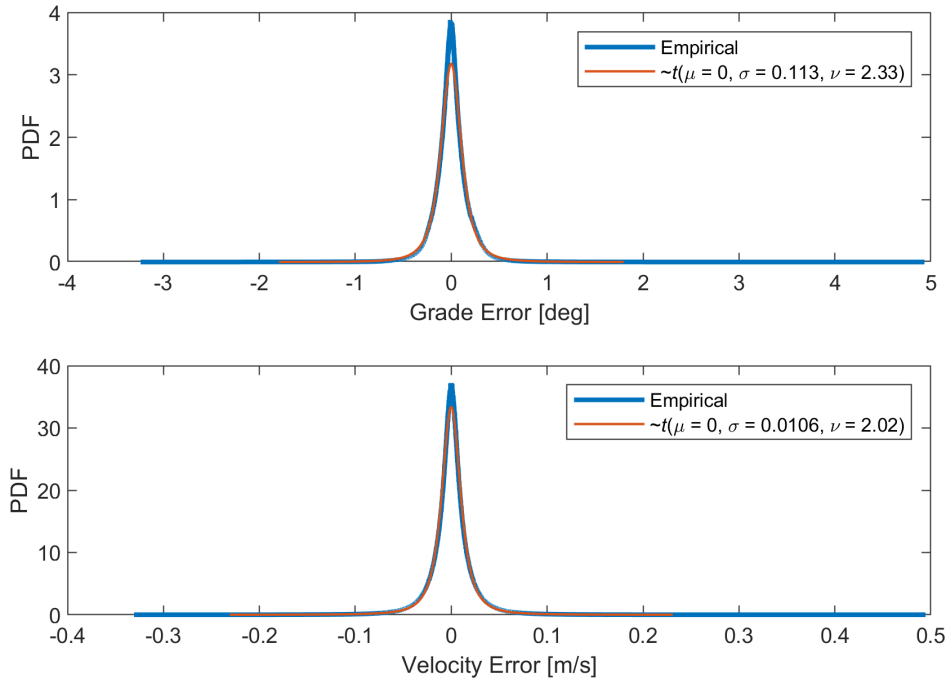

Figure 4.1: Empirical probability density function of the grade and wheelspeed errors, and the accompanying best-fit T-distribution. Data is derived from all four experimental datasets described in the Methodology.

The J1939 CAN torque signal appears to be subject to errors on the order of 2% of the peak torque, although this estimate is an extrapolation from a much different engine [104][2]. Regarding fuel flow, our own analyses have found that CAN-reported fuel flow agrees quite well with the measured fuel flow, although the CAN flow was approximately 6% lower than the measured fuel flow [29]. Without better data on engine efficiency, the power-to-fuel ratio κ can only be estimated.

4.1.2 Sensitivity Analysis Procedure

To perform the sensitivity analysis, both random signal noise and random parameter offsets were introduced to each of the 8000+ IPG TruckMaker simulations (see Section 3.3.1 and Appendix C), one simulation at a time. At the beginning of the analysis, the random number seed is set for reproducibility. First, noise is applied to wheelspeed and grade, as

[2]The cited study has a torque error estimate from an industrial four-cylinder John Deere engine, where the high resolution torque signal was not available. Maximum output and displacement are also much lower than in the trucks in this study.

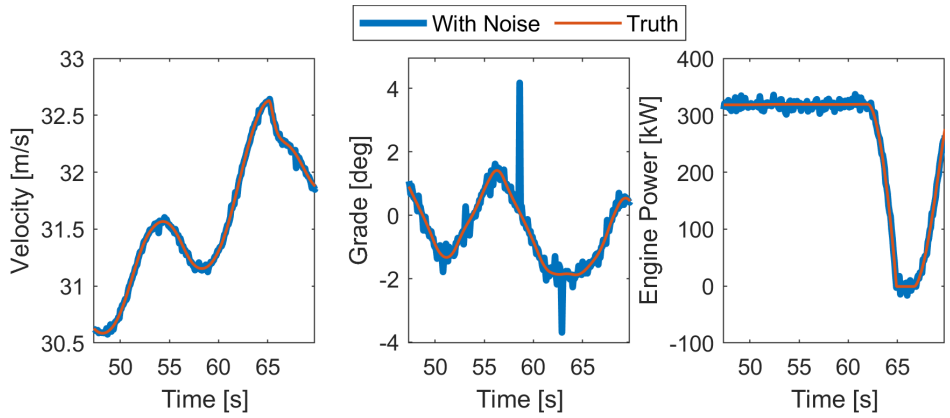

Figure 4.2: The impact of the added noise on the simulated wheelspeed, velocity, and grade.

shown in Figure 4.2. Gaussian error is also added to the engine torque, with a standard deviation of 40 Nm, a rough approximation based on [104]. Then, the parameters C_d, C_{rr}, and m have their offsets chosen randomly from a uniform distribution with a center at zero. A maximum of 50% error was added to C_d and C_{rr}, while for m, an error between $\pm$5000 kg was introduced. It was deemed more realistic to bound the mass error differently than the other parameters, because truck loads are additive, not multiplicative. To clarify, it is reasonable to imagine a fully-loaded truck being overloaded by several tons, but it is less reasonable to imagine that the same truck is 50% heavier than expected. This means that the error in C_d and C_{rr} will come from a pure uniform distribution with sharp edges, but the mass error introduced will come from a mixture of uniform distributions, which has tapered edges. The difference in these distributions can be seen in Figure 4.3, where C_d and C_{rr} error clearly come from a pure uniform distribution, but mass error does not.

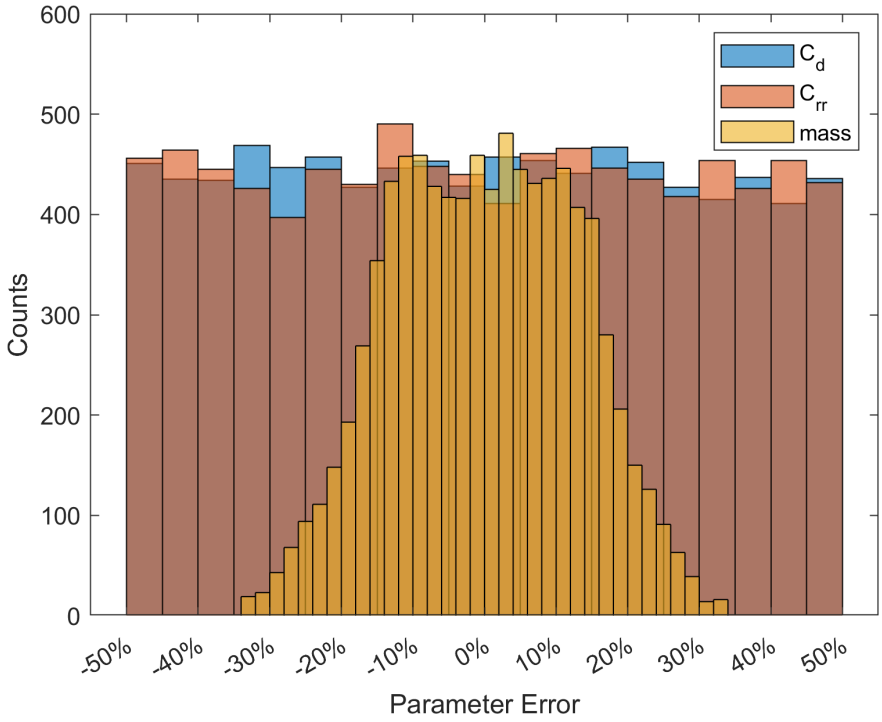

Figure 4.3: Histograms of the three types of parameter errors, expressed as a percentage of the true value.

Once parameter errors have been sampled, then a_{model} and $\hat{a}$ are calculated using noisy signals and offset parameters, allowing the calculation of $NPC_{inferred}$ and its constituent terms. To isolate the influence of each type of parameter error, separate results were generated for drag error, rolling resistance errors, and mass errors, then all three at once.

4.1.3 Sensitivity to C_d Error

Firstly, starting with the aerodynamic drag sensitivity, we shall look at how $NPC_{inferred}$ and all its constituents are affected. An incorrect drag parameter directly affects $\bar{P}_{aero}$. It is not known how an incorrect drag coefficient will affect the estimate of $\bar{P}_{AD}$, nor the combined effect on $NPC_{inferred}$. Recall that $\bar{P}_{AD}$ depends on C_d to generate an intermediate term a_{model}, but that the RLS and Constant Offset procedures attempt to compensate for an incorrect a_{model}.

Figure 4.4 shows the effect of a pure error in C_d on the resulting platoon fuel consumption inference. The percent change in C_d is the independent variable, and the percent change in

inferred values is the dependent variable. As expected, the percent error in $\bar{P}_{aero}$ changes collinearly with the error in C_d. The estimates for $\bar{P}_{AD}$ (one for RLS, the other for Constant Offset) have a very slight negative correlation with the drag error: a 100% change in C_d leads to a change in $\bar{P}_{AD}$ of -1.26% and -0.9% with RLS and Constant Offset, respectively. This is a very small sensitivity. Both methods are effective in compensating for aerodynamic drag errors, and neither RLS nor Constant Offset is clearly superior in this case. We turn

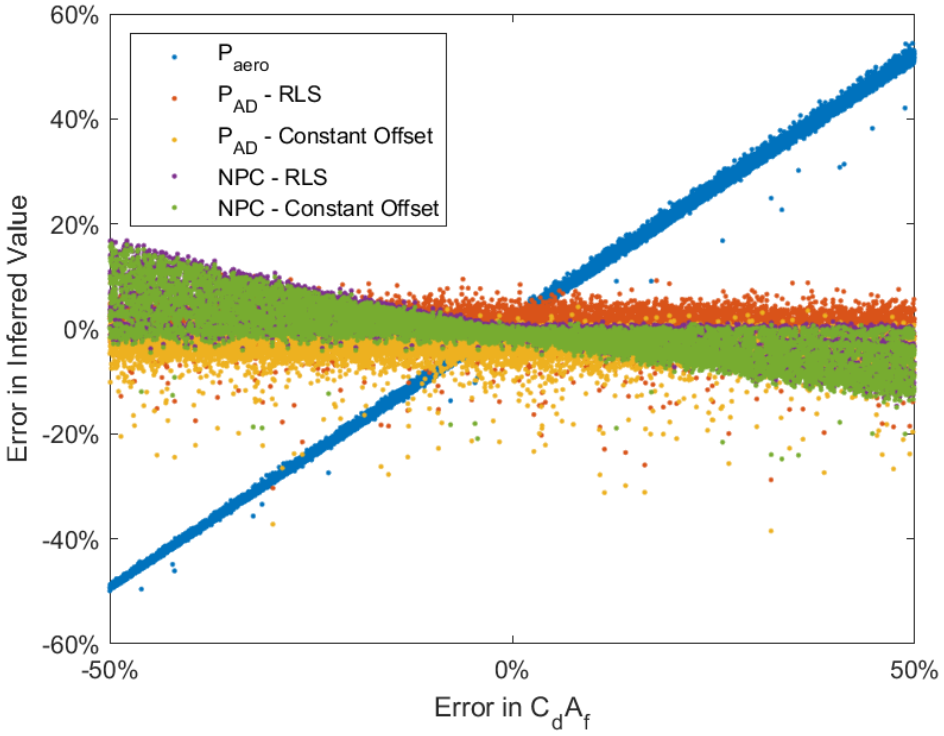

Figure 4.4: The impact of aerodynamic drag error on $NPC_{inferred}$ of a platoon. Parameter errors are uniformly distributed at up to ± 5 m^2, which is 50% of the true value. Percent error is calculated (inferred − true)/true.

our attention to the sensitivity of $NPC_{inferred}$ to aerodynamic drag errors. The shape of the trend is intriguing, spreading out as the aerodynamic drag error increases. When the drag reduction ratio is one, then $\bar{P}_{aero}$ is not of consequence, which explains the cases with large C_d error, yet 0% $NPC_{inferred}$ error. In contrast, the cases most affected by aerodynamic drag modeling errors are those with the greatest drag reduction.

A different way to look at the impact of drag errors is to recreate the $NPC_{inferred}$ versus NPC_{true} plot of Figure 3.16. As illustrated in Figure 4.5, aerodynamic discrepancies

cause a dispersion of $NPC_{inferred}$. Greater displacement of the points occurs at the lower NPC values, resulting in a spear-like visual form. This shape occurs because the higher $NPC_{inferred}$ points are those with less drag reduction, and hence the influence of aerodynamic errors is less significant. We see that the 50% errors in aerodynamic drag shift $NPC_{inferred}$ by much less than 50%.

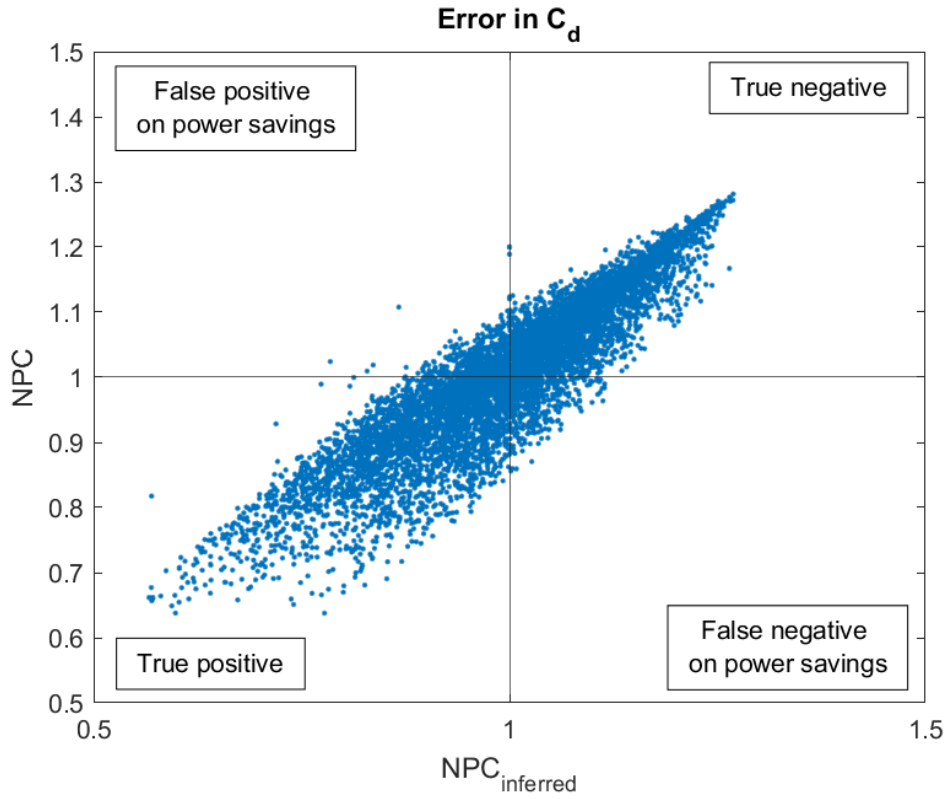

Figure 4.5: The sensitivity of $NPC_{inferred}$ versus true NPC when there is a error in the estimated C_d. $\bar{P}_{AD}$ estimated by RLS method.

One goal of the method proposed in this dissertation is to provide a go/no-go to platoon operators as to whether they saved fuel or not. Figure 4.5 provides a preliminary look at how that may be done. The simplest classification scheme possible would place $NPC_{inferred} < 1$ as go and $NPC_{inferred} > 1$ as no-go. By this scheme, data points in the second and fourth quadrants would be incorrectly categorized.

4.1.4 Sensitivity to C_{rr} Error

With the impact of pure aerodynamic errors covered, we move on to the rolling resistance sensitivity. Since rolling resistance does not play a role in the calculation of $NPC_{inferred}$, we reasonably expect that an error in C_{rr} will not have as strong an effect on the precision of $NPC_{inferred}$ as other errors. As Figure 4.6 shows, this certainly seems to be the case. In addition to a few outliers, the results of $\bar{P}_{AD}$ are within 10% of the true value, and the error of $NPC_{inferred}$ appears to be less than 5%. There is no sensitivity to the percent error in rolling resistance, meaning that $\bar{P}_{AD}$ is being well corrected by both the RLS and the Constant Offset algorithms.

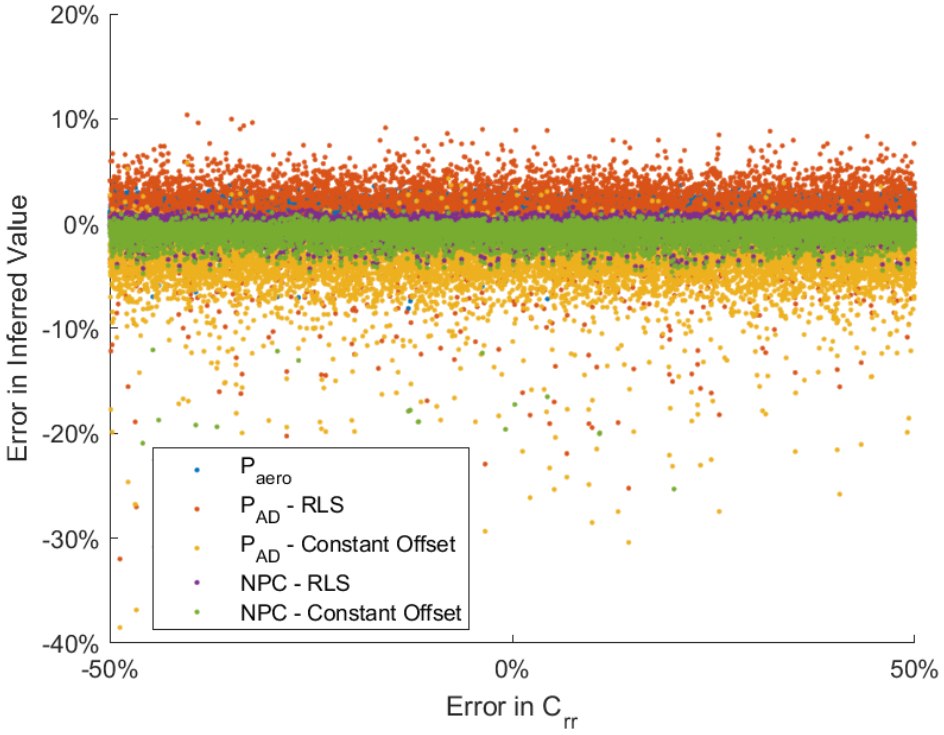

Figure 4.6: The impact of C_{rr} error on $NPC_{inferred}$ and its constituent terms. Parameter errors are uniformly distributed at up to ± 0.005 which is 50% of the true value. Percent error is calculated $\frac{\text{inferred}-\text{true}}{\text{true}}$.

4.1.5 Sensitivity to Mass Error

An error in mass is expected to significantly affect the results of $NPC_{inferred}$. Calculating $\bar{P}_{AD}$ involves multiplying directly by mass; therefore, the inaccuracies in $\bar{P}_{AD}$ are likely to be one-to-one, if not greater. $\bar{P}_{aero}$ is unaffected by inaccurate mass values. As illustrated in Figure 4.7, an increase in mass error leads to an approximately two-fold increase in the error of $\bar{P}_{AD}$. Compared to the Constant Offset algorithm, the RLS algorithm is superior, as indicated by its lower slope (2.37 versus 2.74). The errors in the final estimate of $NPC_{inferred}$ remain mostly below 20% in absolute terms.

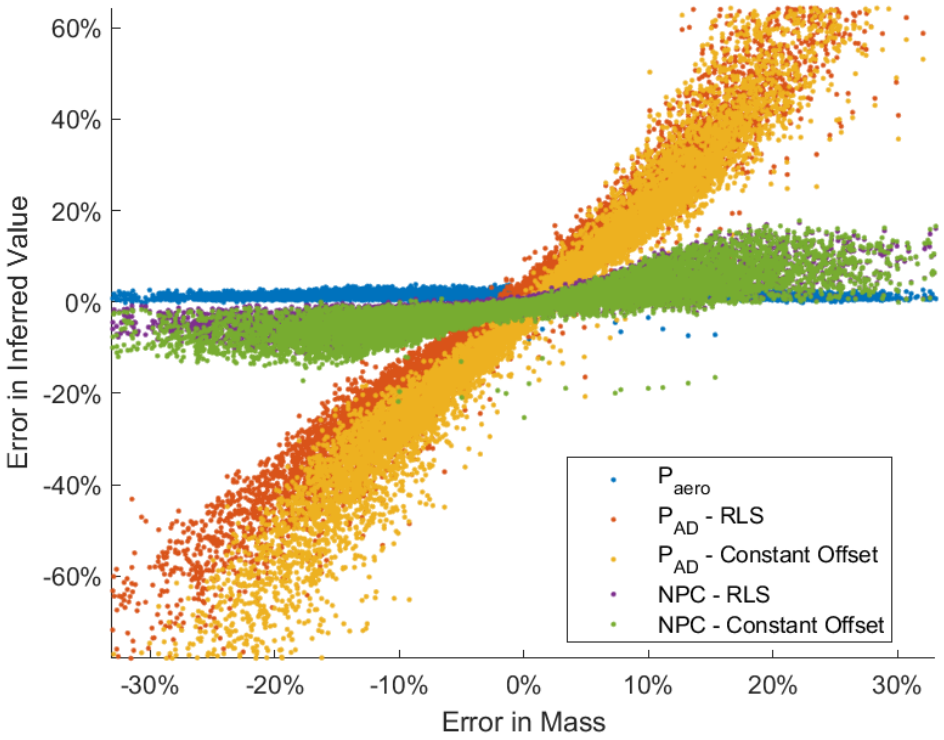

Figure 4.7: The impact of mass error on $NPC_{inferred}$ and its constituent terms. Mass errors are uniformly distributed at up to ±5000 kg. Percent error is calculated $\frac{\text{inferred}-\text{true}}{\text{true}}$.

In essence, an error in mass has the opposite effect as an error in C_d, since it will disproportionately skew the $\bar{P}_{AD}$ estimates. This is evident in Figure 4.8, where it can be seen that the lower $NPC_{inferred}$ values remain relatively unchanged, while higher values, where braking was notable, are scattered.

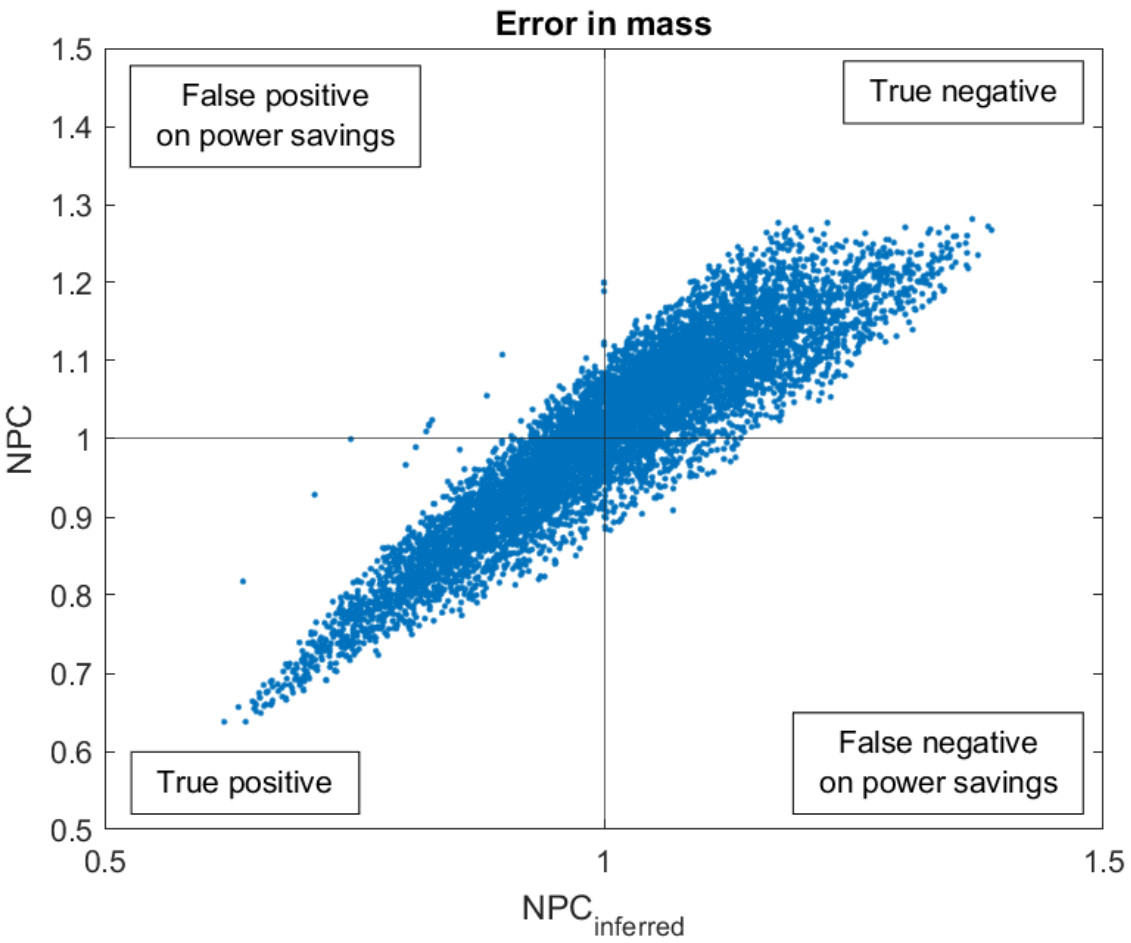

Figure 4.8: The impact of mass errors on inferred power consumption of a platoon. $\bar{P}_{AD}$ estimated by RLS method.

It is worth exploring why there is a greater-than-twofold error in $\bar{P}_{AD}$ when mass errors are introduced. Although an error in mass affect two parts of the modeled acceleration: rolling resistance force and gravitational force. We have already seen that the adaptive $\bar{P}_{AD}$ routines can correct for errors in rolling resistance force. On the other hand, the resistive force of gravity causes a grade-dependent error. Neither of the $\bar{P}_{AD}$ adjustment routines is set up to handle this because they both assume that the error in modeled acceleration will remain constant. Over undulating grade, the actual error in a_{model} will change with the road grade, but the "learned" estimate of error in a_{model} will not. Therefore, it is speculated that grade-dependent acceleration errors are what lead to the greater than twofold error in $\bar{P}_{AD}$.

4.1.6 Combined Sensitivity and Summary

When all three types of parameter errors are combined (along with the noise that was introduced), the net effect is a spread of $NPC_{inferred}$ across the entire simulation range. Figure 4.9 demonstrates this.

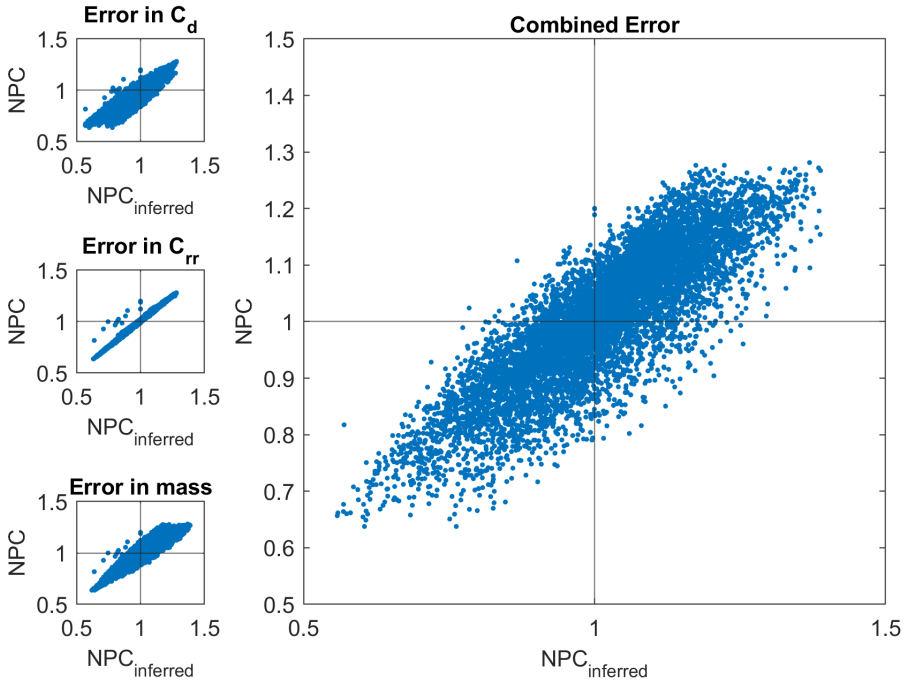

Figure 4.9: The net impact of errors on inferred power consumption of a platoon. Inclusive of noise in velocity, grade, and engine power, and parameter offsets up to ±50% for C_d/C_{rr}, ±5000 kg for mass. $\bar{P}_{AD}$ estimated by RLS.

As a point of reference, 85.8% of the points in Figure 4.9 fall within the true positive and true negative quadrants, meaning that they would have been correctly classified by the inferred power consumption. For further reference, the accuracy with only C_d errors is 88.5%, the accuracy with only C_{rr} errors is 97.8%, and the accuracy with mass errors is 88.9%. It must be noted these accuracy results are specific to the simulation study that was conducted, and that the simulations are not necessarily representative of the population of on-road platooning results.

Based on this sensitivity analysis, the following key conclusions can be drawn:

- Errors in C_d lead to one-to-one error in $\bar{P}_{aero}$ and a very minor impact on $\bar{P}_{AD}$, and affects $NPC_{inferred}$ differently depending on the drag reduction ratio
- Errors in C_{rr} do not significantly affect the accuracy of $NPC_{inferred}$, $\bar{P}_{AD}$, or $\bar{P}_{aero}$
- Errors in mass lead to a two-fold error $\bar{P}_{AD}$
- Error in $\bar{P}_{AD}$ is slightly improved by using the RLS algorithm
- The combined influence of all three parameter errors appears to be additive

Overall, the methodology appears to be relatively robust to the noise and parameter errors that were applied. 50% and 5000 kg are quite large offsets. It should be noted that the noise applied here was stationary and zero mean and that the accessory loads were not modeled. Should a bias or time dependence be introduced, the sensitivity could worsen. Finally, it is time to turn to the experimental results.

4.2 Experimental Results

Now that the sensitivity of our platoon energy prediction scheme has been investigated, we turn our attention to the experimental results with a good idea of the limitations of the prediction. First, we will revisit the Canadian cut-in data from Section 3.2, where existing methods did not accurately predict a 9.7% loss of fuel savings from actuation effort. We will see if $NFC_{inferred}$ offers a better prediction than the previous methods. Looking at the Canadian dataset first also presents a nice example to demonstrate the difference between different drag reduction models and $\bar{P}_{AD}$ adjustment routines.

4.2.1 Canadian Cut-in Revisited

Recall that during one of the tests in Dataset 2, a passenger car inserted itself into the platoon once on each lap. These periodic cut-ins completely erased any fuel benefits for the following vehicle, turning a $10.7 \pm 1.2\%$ benefit into statistically insignificant $1.0 \pm 1.8\%$. The numbers to be matched are:

- A $10.7 \pm 1.2\%$ benefit for the platoon follower at an IVD of 23 meters ($0.881 \leq NPC_{inferred} \leq 0.905$)
- A $1.0 \pm 1.8\%$ benefit for the platoon follower with cut-ins at an IVD of 23 meters ($0.972 \leq NPC_{inferred} \leq 1.008$)
- A difference of $9.7 \pm 2.2\%$ between the pure platooning data and the platooning with cut-ins ($0.881 \leq NPC_{inferred} \leq 0.925$)

If $NPC_{inferred}$ falls within these limits, it is a good sign that the framework works well. We can also check $NFC_{inferred}$, but the value of κ is not well known. Table 4.1 provides all the relevant numbers for calculating $NPC_{inferred}$. The table gives a good idea of all the possible variations on $NPC_{inferred}$ that have been developed in this dissertation. A choice of three braking adjustments, three drag reduction models, with and without weather, leaves 18 different ways to get to $NPC_{inferred}$. Although most variations produce similar values, there are a few surprises, such as the Constant Offset routine, which yields 33% higher $\bar{P}_{AD}$ than without adjustment.

Table 4.1: The relevant values to calculate $NPC_{inferred}$ and $NFC_{inferred}$ for the Canadian cut-in trial. Values in kilowatts and liters/hour.

Sample	$\bar{P}_{AD}$[1]			$\bar{P}_{aero}$[2]		$D\bar{R}R$[3]			$\bar{P}_{total}$	$\bar{Q}_f$
	none	c.ofst	rls	const	wthr	sch	h.rp	h.pw		
Aligned Trial 1	0	0	0	82.55	82.07	0.857	0.805	0.845	104.0	26.68
Aligned Trial 2	0	0	0	82.14	81.17	0.857	0.805	0.845	101.5	26.02
Aligned Trial 3	0	0	0	82.91	81.39	0.857	0.805	0.845	101.5	25.98
Cut-in Trial 1	8.97	11.96	9.01	85.25	83.30	0.857	0.807	0.845	112.1	29.07

[1] "none" is unadjusted, "c.ofst" is Constant Offset, and "rls" is RLS-adjusted
[2] "const" uses constant density and zero wind, and "wthr" includes density differences and estimated wind
[3] "sch" is the Schmid model and "h.rp" and "h.pw" are the Hussein rational polynomial and power models, respectively

Using the consumption results in Table 4.1, we now look at how well $NPC_{inferred}$ agrees with the published fuel savings. Figure 4.10 shows how well each combination of methods for inferring benefits agrees with the actual published savings of the Canadian platoon trials in

[42]. Each of the eighteen combinations of inference methods has been assigned an identifying number from one to eighteen; the interpretation of each number can be found in Table 4.2.

Directing attention first to the top panel of Figure 4.10, it appears that individual predictions for the aligned 23 m IVD case without cut-ins are mostly in agreement with the published savings. Only the ability of the DRR models affects the results in the top panel, since there were no braking losses in the trials. Methods 2, 5, 8, 11, 14, and 17, which all share the Hussein rational polynomial drag reduction model according to Table 4.2, overpredicted platooning benefits.

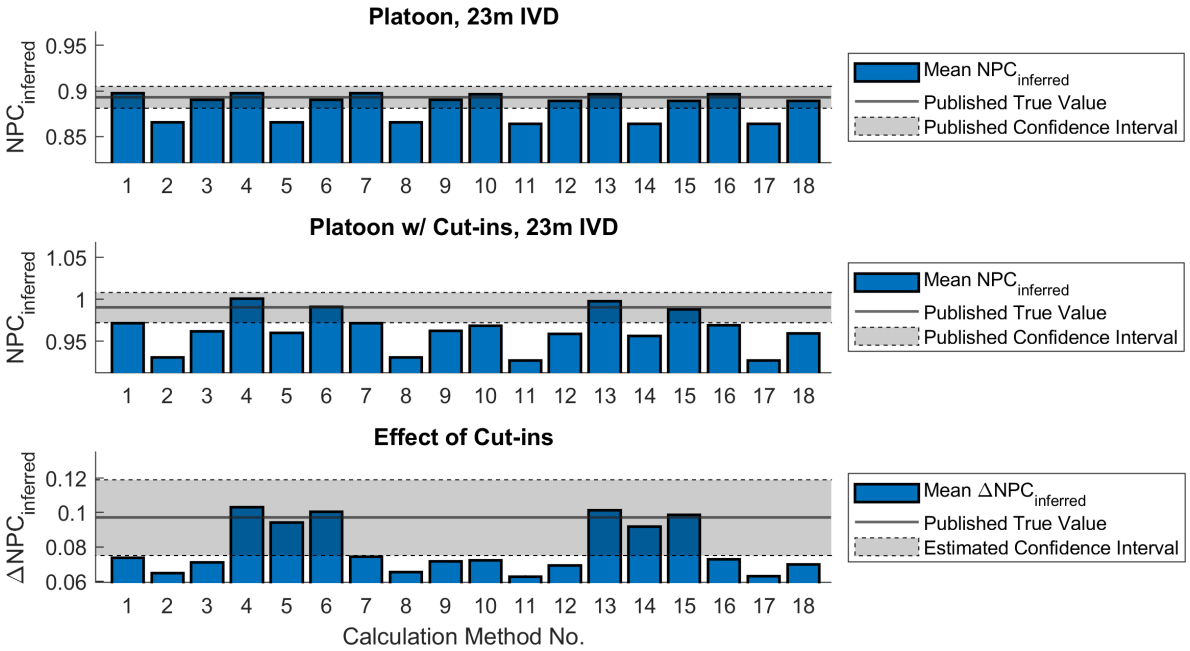

Figure 4.10: $NPC_{inferred}$ for the 23m platoons in [42] without cut-ins (top), with cut-ins (middle), and the difference between the two (bottom). Description of the calculation method numbers is in Table 4.2.

Next, we will focus on platooning with cut-ins, whose results are shown in the middle panel of Figure 4.10. Less of the inferences agree with the published confidence interval in the cut-in case than in pure platooning. This is because for the cut-in case, both aerodynamic benefits and braking losses were present, leaving more potential for error. The methods that agree the most with the published savings all share the Constant Offset PAD adjustment (Methods 4, 6, 13, 15).

Table 4.2: The corresponding methods of the "Calculation Method Numbers" in Figure 4.10.

Method	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
										DRR								
schmid	•			•			•			•			•			•		
h.rp		•			•			•			•			•			•	
h.pw			•			•			•			•			•			•
										$\bar{P}_{AD}$								
No Adjust	•	•	•							•	•	•						
c.ofst				•	•	•							•	•	•			
RLS							•	•	•							•	•	•
										Weather								
with	•	•	•	•	•	•	•	•	•									
without										•	•	•	•	•	•	•	•	•

Finally, we look at the pure impact of cut-ins, which is shown in the bottom panel of Figure 4.10. The pure impact of cut-ins is the difference between the results used to construct the top and middle panels. Because the aerodynamics of the platoons with and without cut-ins were similar, the effect of braking losses dominates the effect of cut-ins as shown in the bottom panel. The six cases that agree are cases 4-6 and 13-15, which are the calculation methods that use the Constant Offset $\bar{P}_{AD}$ adjustment.

In summary, for the majority of the 18 combinations of methods, there is a consistent alignment between $NPC_{inferred}$ and the actual values. It seems that $NPC_{inferred}$ has generally performed well in forecasting platoon energy usage in this specific analysis. However, since this is only a single instance, additional evidence is required before definitive conclusions can be made about the predictive capacity of $NPC_{inferred}$. In theory, we could run a similar analysis comparing the rest of the published fuel benefits from Dataset 2 to the inferred platoon fuel consumption, to see how many of the inferences overlap with the published fuel savings. However, none of the other trials in Dataset 2 featured braking, so this would only show the efficacy of the drag reduction model, not the braking power estimate. Unfortunately, we do not possess a dataset with significant braking and J1321-like accuracy, so such validation will be left to future work.

4.2.2 Consolidating Experimental Data

In this section, a method is devised to consolidate the experimental data from all four datasets into a unified table. Subsequently, this table is used to assess the effectiveness of the platoon energy consumption prediction framework introduced in this thesis. The primary objective in validating the suitability of $NPC_{inferred}$ and $NFC_{inferred}$ as indicators of platooning energy consumption involves computing and contrasting them with the actual NPC_{true} and NFC_{true}. In the same way, we can validate the inferred power/fuel difference ($\Delta P_{inferred}$ and $\Delta F_{inferred}$), though the table will look slightly different.

The minimum viable product is a tool that can provide a useful indication of whether or not a platoon saved fuel. In this case, it is not necessary to express energy consumption as a ratio; a difference works the same way. In some ways, the relative energy benefit is a more satisfying measure than the absolute fuel/power benefit here. With the aim of gamification, operators most likely would want to see their improvement in fuel-distance metrics such as mpg or l/100km. Relative benefit is easier to leverage in this case, as the average fuel economy over the platooning microtrip could be tracked, and the improvement could be displayed to the driver. Fuel savings may be more intuitive to talk about in absolute terms, though over the course of a microtrip, the quantity of fuel saved could be a small, unintuitive number, such as 0.23 Liters of diesel.

No matter whether absolute or relative, we seek to demonstrate the validity of the proposed energy inferences. **We need not show that a platoon saved energy, or did not; we need only show that the energy inferences provided by the framework are highly correlated to the measured energy differences.** To show this correlation, regression analysis will do.

Rather than starting with $NPC_{inferred}$, a more intermediate step would be to compare the difference in drag and active deceleration power in fuel/power between two runs versus the actual energy savings. Similarly to the process used to assess the suitability of previous

methods (Section 3.2), we will regress the drag and braking power onto the observed difference in energy consumption. Specifically, we want to determine whether the drag reduction term correlates with a linear decrease in energy consumption and if the braking term corresponds to a linear increase. **Ideally, the coefficients of the regression parameters would not deviate significantly from one**. This would suggest that the drag reduction and braking estimates are unbiased. Figure 4.11 outlines the process of pairwise comparison that will be employed to test the validity of the framework.

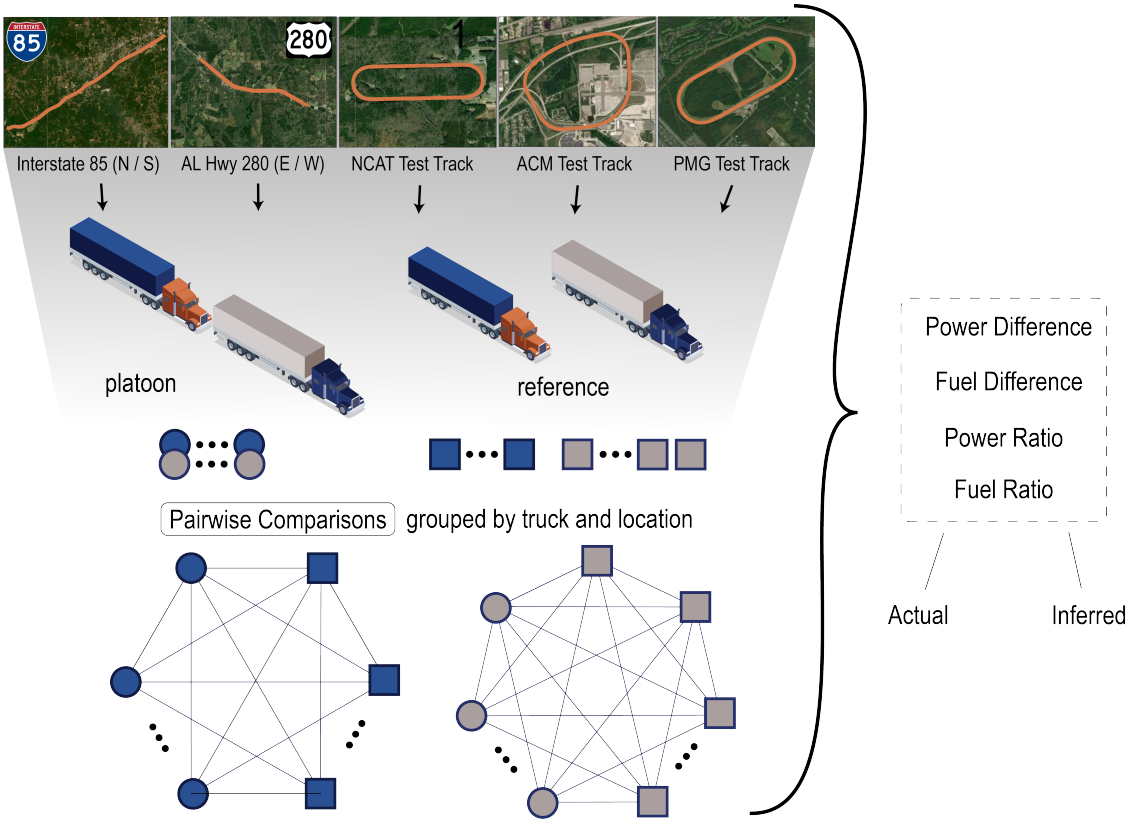

Figure 4.11: The pairwise comparison scheme.

Although it may seem strange, it is not necessary to always compare a platoon run with a standalone run: we could just as easily compare two separate platooning results, as long as the same truck and test segment are used. If we pair up every run from the same truck and road segment, we have arrived at what will be termed the **pairwise model matrix**. The tests will be grouped by truck and test segment (either a test track or a road segment) into

g groups, with n_g tests occurring on a segment, including platoon and baseline runs. Then, for all $\binom{n}{2}$ combinations of tests in each group, the actual and inferred pairwise difference in energy will be calculated.

The pairwise difference structure is allowed under the assumptions that:

- each run is independent
- everything but drag and braking is nominally constant between pairs

The first assumption is reasonably satisfied, since each run occurs at a different time. The second assumption is at least partially satisfied because the pairs are drawn from the same group (same truck, same road segment). Of course, run-to-run differences in weather, accessory loads, and vehicle performance will be present and are sure to increase the variance. To handle these run-to-run differences properly, a mixed-effect model[3] would have to be built, which would be difficult for normalized power and fuel consumption due to their nonlinear structure.

The pairwise model matrix does introduce a reproducibility concern: the regression result will change depending on the order of the combinations, since differences and ratios depend on order. To avoid introducing a dependency structure in the output table due to the order of combinations, the selections will be flipped randomly[4].

To sum up the preceding discussion of how the validation will be conducted, we will show the construction of the pairwise model matrix in Tables 4.3 through 4.6. Starting with the source data, Table 4.3 is constructed with one replicate per row, each indexed by group G and the in-group index N. Some of the datasets have control truck data paired with each replicate, and other do not. To encode all datasets in the pairwise matrix equally, we will set the control consumption to a constant value (1) when no control truck was available for a replicate.

[3]More specifically, a multiple-membership model would have to be built, which is a subset of mixed-effect models.

[4]This reproducibility concern could also be addressed by taking the absolute value/reciprocal such the result is always of a given sign or above 1, but this may create more issues than it solves.

Table 4.3: The initial model matrix before data is randomly paired.

G *(groups)*	N_g *(in group)*	T *(test P)*	C *(control P)*	P_{AD} *(braking P)*	P_{aero} *(aero P)*	DRR *(drag reduction)*
1	1	$T_{1,1}$	$C_{1,1}$	$P_{AD\ 1,1}$	$P_{aero\ 1,1}$	$DRR_{1,1}$
⋮	⋮					
1	n_1					
2	1					
⋮	⋮					
g	n_g	T_{g,n_g}	$\cdots$			

Then, Table 4.3 is augmented to include all possible pairings for each group G, as illustrated in Table 4.4. The initial sample in every pair is denoted by the subscript *plat*, while the second is marked with the subscript *ref*. However, it is arbitrary whether or not a sample originates from a platoon. Both samples can be from a platoon, or one, or neither, as long as they are distinct samples from the same truck on the same route.

Table 4.4: The model matrix after data is randomly paired.

G	N_{plat}	N_{ref}	T_{plat}	T_{ref}	C_{plat}	C_{ref}	$P_{AD,\,plat}$	$P_{AD,\,ref}$	$P_{aero,\,plat}$	$P_{aero,\,ref}$
1	1	2	$T_{1,1}$	$T_{1,2}$	$C_{1,1}$	$C_{1,2}$	$P_{AD\ 1,1}$	$P_{AD\ 1,2}$	$DRR\times P_{aero\ 1,1}$	$DRR\times P_{aero\ 1,2}$
⋮	⋮	⋮								
1	$\binom{n_1}{2}$									
2	1	2								
⋮	⋮	⋮								
g	$\binom{n_g}{2}$		$\cdots$							

Finally, applying numerous transformations to the variables in Table 4.4, we arrive at the final pairwise model matrices, as shown in Table 4.5 and 4.6. Table 4.5 will be used to assess the precision of relative energy change as a ratio ($NPC_{inferred}$), while Table 4.6 provides the pairwise difference in absolute terms (with units of kW or L/hr). The notation $\Delta(P_{AD} + P_{aero})_{plat-ref}$ is short for $(P_{AD\ plat} + P_{aero\ plat}) - (P_{AD\ ref} + P_{aero\ ref})$.

Note that in Table 4.6, ΔP_{true} is multiplied by the average control consumption of the pair, to return the difference in the T/C ratios to absolute terms. Also, the terms ΔP_{AD}

and ΔP_{aero} have been retained, because it will be useful more insightful to apply multiple linear regression to each one rather than just fitting $\Delta P_{inferred}$ versus ΔP_{true}.

Table 4.5: The final model matrix for assessing relative platooning energy consumption.

G	N_{plat}	N_{ref}	NPC_{true}	$NPC_{inferred}$
1	1	2	$\left[\frac{T_{plat}/C_{plat}}{T_{ref}/C_{ref}}\right]$	$\frac{T_{plat}}{T_{plat}+\Delta(P_{AD}+P_{aero})_{plat-ref}}$
$\vdots$	$\vdots$	$\vdots$		
1	$\binom{n_1}{2}$			
2	1	2		
$\vdots$	$\vdots$	$\vdots$		
g	$\binom{n_g}{2}$		$\cdots$	

Table 4.6: The final model matrix for assessing absolute platooning energy consumption.

G	N_{plat}	N_{ref}	ΔP_{true}	ΔP_{AD}	ΔP_{aero}	$\Delta P_{inferred}$
1	1	2	$\left[\frac{T_{plat}}{C_{plat}}-\frac{T_{ref}}{C_{ref}}\right]\times\frac{C_{plat}+C_{ref}}{2}$	$P_{AD,\,plat}$ $-P_{AD,\,ref}$	$P_{aero,\,plat}$ $-P_{aero,\,ref}$	$\Delta P_{AD}+\Delta P_{aero}$
$\vdots$	$\vdots$	$\vdots$				
1	$\binom{n_1}{2}$					
2	1	2				
$\vdots$	$\vdots$	$\vdots$				
g	$\binom{n_g}{2}$		$\cdots$			

To rewrite the matrices in terms of fuel, one substitutes fuel for power on the true values and multiplies the inferred power values by the power-to-fuel conversion factor κ, as was done in Equation 3.24. The actual regression table that will be used to create the results will have more columns than the Tables 4.6 and 4.5, because of the redundant drag reduction and braking power estimates from different drag models and braking estimates.

All individual trials from each dataset can be placed in Tables 4.5 and 4.6. In doing so, the entire set of available platooning data can be used at one time, despite the differing experimental designs. This will greatly strengthen the validation, since some datasets had very little difference in aerodynamic drag savings, and others had very little braking.

The usefulness of $NPC_{inferred}$ is reinforced by looking at the Tables 4.5 and 4.6. There are many situations in which the reference $P_{AD,\,ref}$ or $P_{aero,\,ref}$ can be considered zero. Under this condition, platoon energy consumption can be directly inferred in real-time, without the need for standalone truck data. Conversely, the true measures absolutely require a standalone references; there is no way around it. Whether the inference matches the actual result is another question altogether and will be investigated in short order.

4.2.3 Motivation for Robust Regression

Due to the pairwise model matrix, outliers will appear several times each, since observations are members of several pairs. A multiple-membership model would be required to fully handle the pairwise structure [116], although specifying and interpreting such a model is beyond the scope of this dissertation. Instead, we turned to robust regression to lessen the influence of outliers and get more consistent estimates of the coefficients of interest. The robust regression that will be employed is the MATLAB implementation of Tukey's bisquare method, which uses iteratively weighted least squares to assign weights to each observation. Tukey's bisquare method does allow observations to be down-weighted to zero if their residual is large enough.

4.2.4 Representative Result

In this subsection, a representative result is shown that includes all the experimental data, first in absolute terms, then in relative terms. The representative result shown here uses the Schmid DRR model, Constant Offset of P_{AD}, and no weather lookup. As explained previously, a robust regression will be used to reduce the influence of outliers on the fitted coefficients.

The results of this subsection were obtained by only one of the 18 combinations of methods for calculating the energy inferences, so there are 17 additional combinations of methods available for analysis. This dissertation proposes a framework for the prediction of

platooning energy benefits, and we endeavor to show its inner workings in this section. Doing so will help to understand results going forward, as well as provide a methodical approach to analyzing any future drag reduction and braking loss models.

Absolute Power Difference

The relative inference for the change in platooning energy ($NPC_{inferred}$) is inclusive of our absolute measure ($\Delta P_{inferred}$) so it is wise to start with absolute terms first. Furthermore, $\Delta P_{inferred}$ is itself composed of the difference in braking and drag power (ΔP_{AD} and ΔP_{aero}). To start, we can make use of multiple regression techniques on the model matrix of Table 4.6 by separately fitting the true power difference (ΔP_{true}) versus ΔP_{AD} and ΔP_{aero}.

We wish to assess whether the difference in braking power and drag power is a one-to-one relationship with the difference in actual power consumption. In a linear modeling sense:

$$\Delta P_{true} = \beta_0 + \beta_1 \Delta P_{AD} + \beta_2 \Delta P_{aero} + \varepsilon \tag{4.1}$$

$$\beta_1 \stackrel{?}{=} \beta_2 \stackrel{?}{=} 1 \tag{4.2}$$

If Equation 4.2 is approximately satisfied, then it is a good indication that we can predict absolute platooning energy change.

Figure 4.12 shows a three-dimensional graph of the absolute power difference versus the drag and braking difference. Amazingly, all four datasets agree quite well. The R^2 of the regression plane is 0.677, which is a moderate correlation. The y-intercept of the response plane shown in Figure 4.12 is effectively zero.

The one-to-one plane that satisfies equation 4.2 is also shown in Figure 4.12. It is clear that the fitted plane is not much different from the one-to-one plane, both numerically and visually. The coefficient for ΔP_{AD} is close to one (0.98 ± 0.01), but the coefficient for ΔP_{aero}

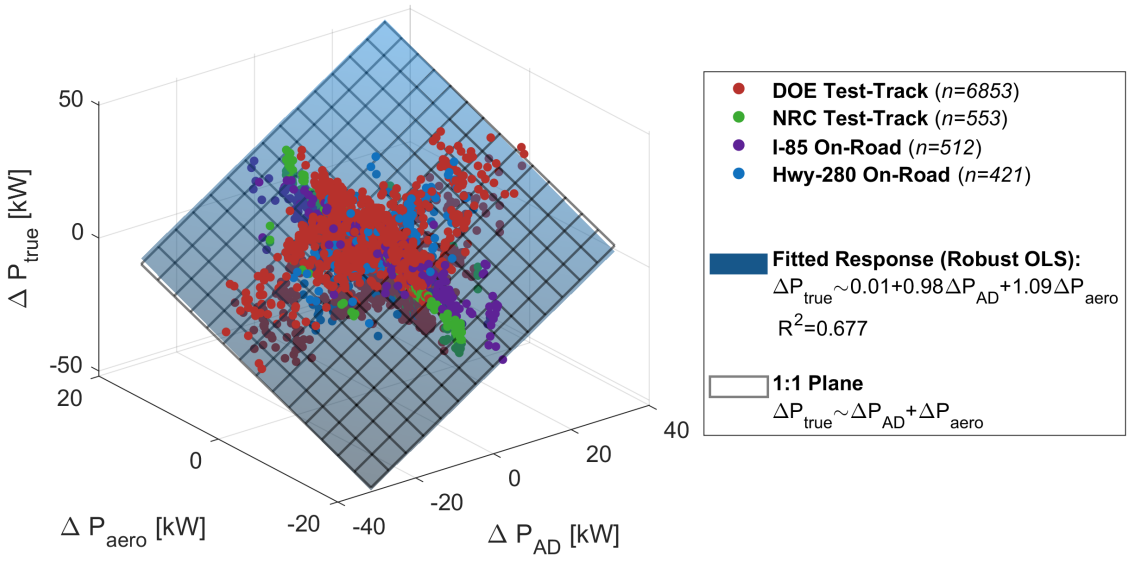

Figure 4.12: The fitted regression plane of true power difference versus the braking and drag power difference.

is higher than one (1.09 ± 0.01). This means that the model underestimated the aerodynamic benefits on average, leading to conservative estimates of energy savings.

So much is plotted in the same space in Figure 4.12 that it is difficult to tell how well each dataset matches the one-to-one gridded plane. We can alter the regression equation to include a dummy variable $\mathbf{D}$ for the dataset, which allows the slopes to change for each dataset.

$$\Delta P_{true} = \beta_0 + \beta_{\mathbf{1,D}} \Delta P_{AD} \mathbf{D} + \beta_{\mathbf{2,D}} \Delta P_{aero} \mathbf{D} + \varepsilon \tag{4.3}$$

The notation in Equation 4.3 is compact; in reality, there will be a pair of slope coefficients for each dataset. In our case, there are four datasets, so there will be eight different slope coefficients.

Figure 4.13 shows how the regression plane deviates when it is allowed to change for each dataset. The unique aspects of each dataset are clearly on display in Figure 4.13. For example, there was not much braking in Dataset 2 or 3, so the fit is uncertain with respect to ΔP_{AD}. Another observation: Dataset 3 has a good distribution of drag reduction, but the ΔP_{aero} slope is significantly less than one. One possible explanation is that the dataset was

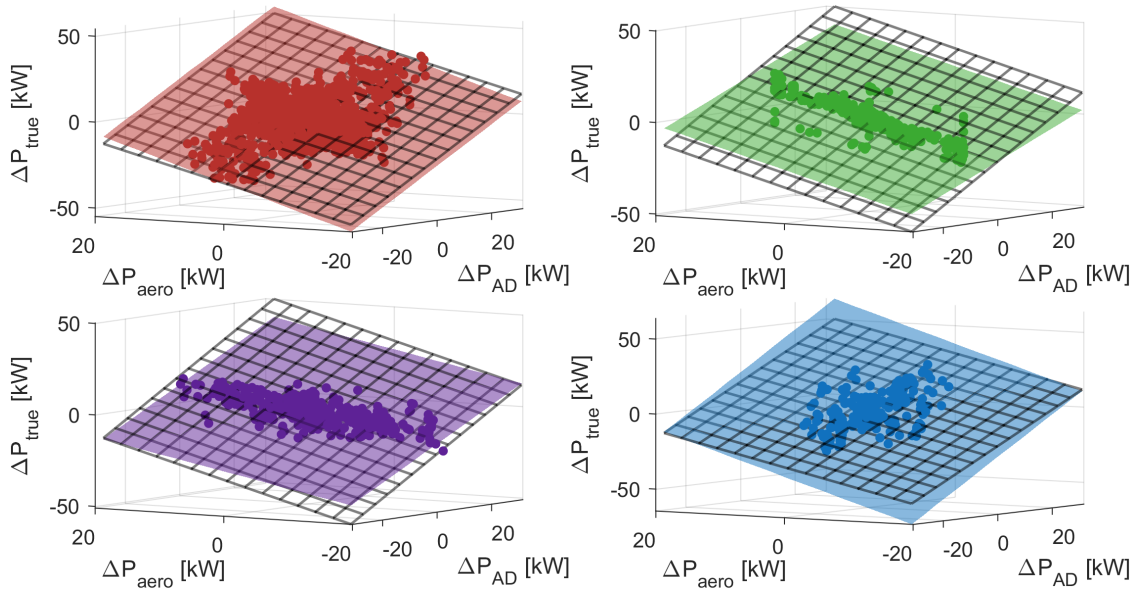

Figure 4.13: The fitted regressions planes for each dataset of absolute power difference versus braking and drag differences. The gridded plane in each panel is the one-to-one plane. The datasets are: Dataset 1: DOE test-track (red, upper left), Dataset 2: NRC (green, upper right), Dataset 3: I-85 (purple, lower left), Dataset 4: Hwy-280 (blue, lower right).

conducted in the presence of surrounding traffic, which could lead to an overestimation of aerodynamic benefits. The slopes shown in Figure 4.13 are quantified in Table 4.7. Among

Table 4.7: The slopes of the planes in Figure 4.13 for inferences on absolute power difference. Slopes which are no different than one are in bold.

	$\beta_{\mathbf{1,D}}$ (affects ΔP_{AD})	$\beta_{\mathbf{2,D}}$ (affects ΔP_{aero})
Dataset 1	$\mathbf{0.9976} \pm 0.0197$	1.2132 ± 0.0252
Dataset 2	0.6915 ± 0.1141	$\mathbf{0.9931} \pm 0.0414$
Dataset 3	0.8499 ± 0.1452	0.7190 ± 0.0546
Dataset 4	1.2028 ± 0.0579	1.3742 ± 0.1529

the eight slopes, two are not statistically different from one. For the slopes that are not equal to one, the following comments are made:

- The coefficient of ΔP_{AD} for Datasets 2 and 3 is lower than 1. There is scarce data in the braking axis for these datasets, which could lead to an incorrect slope.
- The coefficient of ΔP_{AD} for Datasets 4 is higher than 1. There was much braking in this dataset, so there is no easy explanation for this lack of fit.

- The coefficient of ΔP_{aero} for Dataset 1 is much higher than one (1.22 ± 0.03). This dataset includes many heterogeneous four-truck platoons whose aerodynamic benefits are not well understood.
- The coefficient of ΔP_{aero} for Dataset 3 is lower than one (0.74 ± 0.07). This is partially expected due to surrounding traffic, which reduces the potential for aerodynamic benefits.
- The coefficient of ΔP_{aero} for Dataset 4 is greater than one (1.24 ± 0.18). There was only one headway for the entire dataset, so the data in the drag axis is relatively scarce.

Absolute Fuel Results

Having investigated the absolute power results, we now will look at the absolute fuel results. The regression equation for the absolute fuel difference is nearly identical to the absolute power difference equation (Equation 4.1), but with the addition of a power-to-fuel ratio, κ, which depends on the fuel properties and engine efficiency η_{gen} (Equation 3.22).

$$\Delta F_{true} = \beta_0 + \kappa(\beta_1 \Delta P_{AD} + \beta_2 \Delta P_{aero}) + \varepsilon \tag{4.4}$$

$$\beta_1 \stackrel{?}{=} \beta_2 \stackrel{?}{=} 1 \tag{4.5}$$

An initial estimate of κ is made assuming that the diesel fuel has 36 MJ/L and the engine is 36.6% efficient based on comparison between CAN-reported power and fuel in [10]. However, this results in a plane with both slope coefficients greater than one. The slope for braking loss is 1.19 and the slope for drag reduction is 1.08. We outline two possibilities, which are not exclusive:

1. The true value of η_{gen} and/or the fuel energy density are significantly offset from the estimated value
2. The powertrain efficiency is lowered (or otherwise affected) by platooning dynamics

Both of these explanations are plausible, given the simplification of assuming a constant engine efficiency. An η_{gen} of 32.2% places both slopes as close to one as possible, as shown in Figure 4.14. Perhaps the tank-to-wheel efficiency of the engines is lower than previously suspected.

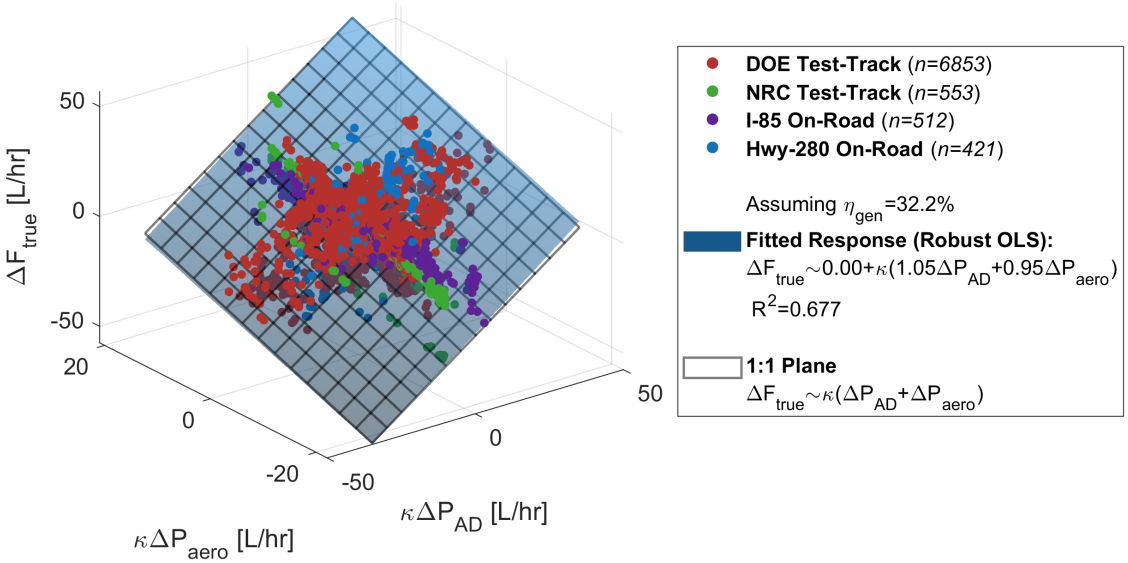

Figure 4.14: The fitted regression plane of true fuel rate difference versus the inferred fuel difference due to braking and drag reduction. $\kappa = (0.322 \cdot 36e6)^{-1}$ liters per Joule.

If we allow the slopes to vary with each dataset as was done for the absolute power difference, it allows us to assess the assumption that normalized power consumption is approximately equal to normalized fuel consumption. Table 4.8 gives the slopes of the regression planes when using absolute fuel (using $\eta_{gen} = 32.2\%$). It is the twin of Table 4.7, so the percent change of each slope (from Table 4.7 to Table 4.8) has also been listed.

Table 4.8: The slopes of absolute fuel difference versus inferred fuel difference due to braking and drag differences. Slopes which are no different than one are in bold. The percentages in italics are the percent change of the slope versus its corresponding value in Table 4.7, which is the power equivalent of this table.

$\eta_{gen} = 32.2\%$	$\beta_{\mathbf{1,D}}$ (affects $\kappa\Delta P_{AD}$)		$\beta_{\mathbf{2,D}}$ (affects $\kappa\Delta P_{aero}$)	
Dataset 1	1.0739 ± 0.0190	(*7.65%*)	1.1096 ± 0.0244	(*-8.54%*)
Dataset 2	0.6475 ± 0.1103	(*-6.35%*)	0.8172 ± 0.0400	(*-17.71%*)
Dataset 3	0.7133 ± 0.1404	(*-16.08%*)	0.5815 ± 0.0528	(*-19.13%*)
Dataset 4	1.8530 ± 0.0560	(*54.05%*)	1.1691 ± 0.1478	(*-14.92%*)

Most slopes are very different when expressed in fuel terms rather than in power terms. Most significantly, the slope related to braking for Dataset 4 is 54% higher. If the regression can be trusted, this means that the fuel cost of braking for this dataset was 54% higher than the power cost. This could happen if braking induced a large number of gear changes and the efficiency of the powertrain was drastically reduced, which is a violation of the assumption of constant powertrain efficiency (Assumption 2). However, if different η_{gen} were assumed, then the slopes also change, so we should be careful not to overanalyze.

Relative Power and Fuel Results

In general, our methods predicted the absolute difference in vehicle energy consumption well, and we can expect a similar result for relative results. These results are given in Figure 4.15, in terms of power in Figure 4.15a, and fuel in Figure 4.15b. For power (4.15a), the slope of the fit line is 1.018, which is very close to one and the R^2 value is 0.662. There are numerous outliers that scatter vertically, which is in part due to the pairwise model matrix.

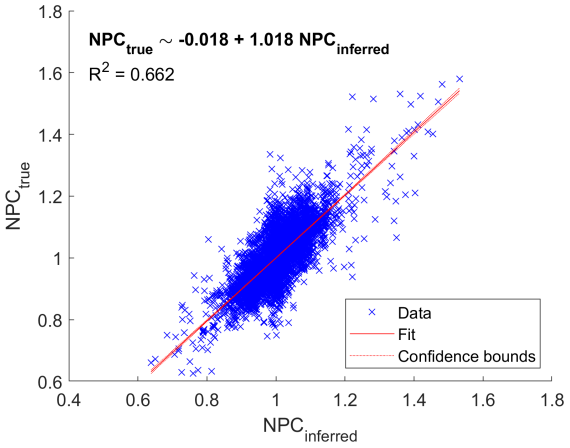

(a) NPC_{true} versus $NPC_{inferred}$ for the all combined experimental datasets.

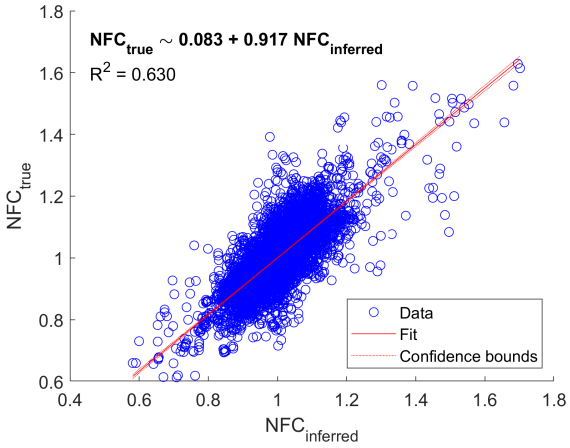

(b) NFC_{true} versus $NFC_{inferred}$ for the all combined experimental datasets, assuming $\eta_{gen} = 32.2\%$.

Figure 4.15: A relative platoon energy consumption result, in terms of power (left) and fuel (right).

The relationship of NFC versus $NFC_{inferred}$ in Figure 4.15b is nearly identical to the NPC plot. The only difference is that the slope is significantly lower than one (0.914). As

for why the slope is less than one, this is not so simple. The sources of discrepancy between inference and true values are easily explained when working in absolute terms, but in relative terms, many factors are lumped together.

Summary

In this subsection, a representative result is provided for inferences of absolute and relative platooning energy consumption, using the Schmid drag reduction model, Constant Offset adjustment, and constant weather. The absolute power difference between each pair of comparable runs (ΔP_{true}) was fairly well captured by the differences in drag reduction and braking for those runs (ΔP_{AD} and ΔP_{aero}). Absolute fuel results indicate that the inferred change in fuel consumption is less than the actual change. This could be the by-product of an incorrect κ or a real physical phenomenon that decreases the powertrain efficiency during control-induced accelerations (e.g. gear-shifts); we do not possess the tools to tell which. The chosen value for κ used an LHV of $36{\times}10^6$ MJ/L and an engine efficiency of $\eta_{gen} = 32.2\%$. Relative results in the forms of $NPC_{inferred}$ and $NFC_{inferred}$ are also shown. These results agree with the absolute results, but do not provide as much specific information about the adequacy of the fit.

4.2.5 All Results

Having taken an in-depth look at one of the results, let us now expand our horizons to include all of the different methodologies. We seek to establish the limitations of our estimated platooning benefits.

Table 4.9 presents the coefficient of determination (R^2) for the robust least squares regression of NPC_{true} on $NPC_{inferred}$ across 18 different combinations of drag reduction modeling, braking power adjustment, and weather fidelity.

The following conclusions can be drawn from the table:

Table 4.9: The R^2 of each method on $NPC_{true} \sim NPC_{inferred}$. Asterisk indicates the highest R^2 in row or column, and the **bold** values are the average for each 3-by-3 subtable.

Weather	Constant Weather				Queried Weather			
$\bar{P}_{AD}/\bar{P}_{aero}$	None	RLS	C.Offset	*Avg.*	None	RLS	C.Offset	*Avg.*
Schmid	0.697	0.675	0.681	0.684*	0.663	0.660	0.661	0.661*
Hussein RP	0.682	0.648	0.659	0.663	0.673	0.652	0.659	0.661*
Hussein Pwr	0.663	0.629	0.639	0.644	0.634	0.614	0.621	0.623
Avg.	0.681*	0.651	0.660	**0.664**	0.657*	0.642	0.647	**0.649**

- **Regarding $\bar{P}_{AD}$:** surprisingly, no P_{AD} adjustment yielded the highest R^2, followed by Constant Offset, then RLS.
- **Regarding $\bar{P}_{aero}$:** The three drag reduction models are similarly correlated with only minor differences. In general, the Schmid model does the best, though it is the only one that explicitly allows four-truck platoons.
- **Regarding weather:** The way that queried air density and headwind were incorporated into the framework did not lead to a better correlation according to R^2. This does not mean that wind and weather are not important; their effect should be reevaluated with a more accurate model.

Similar trends were observed when Table 4.9 was constructed using ΔP_{true}, ΔF_{true}, and NFC_{true} instead of NPC_{true}. Because the rough headwind approximation of weather effects did not add explanatory power, it will be omitted from the upcoming tables to make room for fuel results.

R^2 is far from the only important metric. We also want to inspect the slopes of the regression to assess the precision of our inferences on platooning energy consumption. Absolute results are the best place to start, due to the ability to separate the coefficients for $\bar{P}_{AD}$ and $\bar{P}_{aero}$.

First, we turn attention to the coefficients associated with braking differences (ΔP_{AD}), which indicates the extent to which the braking power is represented. In the sensitivity analysis discussed earlier in this chapter (Section 4.1), it was observed that the RLS and

Constant Offset procedures effectively addressed errors in vehicle parameters, but they were unable to eliminate the impact of mass inaccuracies. If this is the case for the experimental data, both methods should yield coefficients closer to 100%.

Table 4.10 shows the resulting coefficient related to ΔP_{AD}, when a regression on the power difference is performed. An optimistic braking coefficient would be less than 100%, whereas a pessimistic one would be greater than 100%, because braking losses are an undesirable aspect of platooning. Overall, the Constant Offset method demonstrates the most accuracy in capturing actual braking losses (closest to 100%), and unadjusted $\bar{P}_{AD}$ is the most precise (lowest standard error). However, unadjusted $\bar{P}_{AD}$ yields pessimistic predictions of the braking losses on average, in the 85% range for power. Adjustment routines yield coefficients closer to 100% than the unadjusted estimate of braking losses in terms of power. The drag reduction model does affect the braking estimate by one or two percentage points, but all of the intervals overlap in each column.

The fuel slopes are subject to the specified κ, which would scale the percentages in the right half of Table 4.10 accordingly. Here, κ has been set to $0.322 \cdot 36 \times 10^6$ liters per Joule, representing a 32.2% power generation efficiency and a fuel energy density of 36 MJ/L. This value was chosen based on the representative results, but it is not firmly justified by any stretch.

Table 4.10: The slope of the coefficient related to $\bar{P}_{AD}$ (β_1), expressed as a percentage of the true value. For fuel results, η_{gen} was set to 32.2%. Values in **bold** are those that are not significantly different from 100%, according to the 95% confidence interval.

Regression	$\Delta P_{true} \sim \beta_0 + \beta_1 \Delta P_{AD} + \beta_2 \Delta P_{aero}$			$\Delta F_{true} \sim \beta_0 + \kappa(\beta_1 \Delta P_{AD} + \beta_2 \Delta P_{aero})$		
$\bar{P}_{AD}/\bar{P}_{aero}$	None	RLS	C.Offset	None	RLS	C.Offset
Schmid	87.8 ± 1.6	103.8 ± 2.0	$\mathbf{98.2} \pm 1.8$	$\mathbf{98.6} \pm 1.5$	108.9 ± 1.9	104.7 ± 1.8
Hussein RP	86.7 ± 1.6	102.6 ± 2.0	96.9 ± 1.9	97.0 ± 1.6	107.8 ± 2.0	103.4 ± 1.8
Hussein Pwr	86.1 ± 1.7	$\mathbf{101.7} \pm 2.1$	96.4 ± 1.9	96.2 ± 1.6	107.3 ± 2.0	103.4 ± 1.9

Having looked at the coefficient associated with braking, we next inspect the coefficient corresponding to aerodynamic drag reduction ($\bar{P}_{aero}$). Suppose that the assumed air density of 1.225 kg/m^3 is reflective of the actual environmental conditions that the trucks experienced

and that the drag coefficients and front areas of the trucks are correctly specified. If so, then the coefficient associated with $\bar{P}_{aero}$ reveals the bias inherent in each drag reduction model. Similarly to the methods used to estimate $\bar{P}_{AD}$, the estimate of $\bar{P}_{aero}$ can be overly optimistic or pessimistic with respect to drag savings. However, the interpretation differs, as drag savings are a favorable aspect of platooning. A coefficient that exceeds 100% indicates a pessimistic view of drag savings, while a value below 100% suggests optimism.

As indicated in Table 4.11, only the Schmid model demonstrates pessimism towards drag savings (approximately 110%). In order of increasing optimism are the Hussein Power model (approximately 95%) and the Hussein Rational Polynomial model, which exhibits the highest level of optimism (approximately 90%).[5] For the fuel results in the right half of Table 4.11, the coefficients scale with κ. At the chosen κ, drag reduction appears to affect fuel less than it affects power, since the coefficients are lower for fuel than they are for power. A higher value of η_{gen} would lead to higher fuel coefficients.

Table 4.11: The slope of the coefficient related to $\bar{P}_{aero}$ (β_2), expressed as a percentage of the true value, with constant air density and no winds. For fuel results, η_{gen} was set to 32.2%. Values in **bold** are those that are not significantly different from 100%, according to the 95% confidence interval.

Regression	$\Delta P_{true} \sim \beta_0 + \beta_1 \Delta P_{AD} + \beta_2 \Delta P_{aero}$			$\Delta F_{true} \sim \beta_0 + \kappa(\beta_1 \Delta P_{AD} + \beta_2 \Delta P_{aero})$		
$\bar{P}_{AD}/\bar{P}_{aero}$	None	RLS	C.Offset	None	RLS	C.Offset
Schmid	109.8 ± 1.9	109.6 ± 2.0	108.6 ± 2.0	**98.2** ± 1.9	96.1 ± 1.9	95.3 ± 1.9
Hussein RP	90.9 ± 1.7	90.5 ± 1.7	89.7 ± 1.7	80.3 ± 1.6	78.5 ± 1.7	77.9 ± 1.7
Hussein Pwr	94.6 ± 1.9	95.0 ± 1.9	94.1 ± 1.9	84.5 ± 1.8	83.1 ± 1.9	82.5 ± 1.8

The slopes provided for braking losses and drag reduction are diagnostic in nature, helping us to understand the comprehensive results when the two are summed together. It is not the goal to empirically adjust $\bar{P}_{AD}$ and $\bar{P}_{aero}$ based on the results, but rather to see how

[5]It is important to note that the Hussein models were not originally designed to accommodate four-truck platoons (they were designed for a maximum of three trucks). Therefore, the drag reductions for the four-truck configurations were computed by treating the middle trucks equally and considering the fourth truck as the third. This assumption would presumably lead to more pessimistic results for the Hussein models, making it intriguing that they demonstrate optimism towards drag savings.

close the inferences come to the truth. So next, we will look at the total inferred difference $\Delta P_{inferred}$, which is the sum of ΔP_{AD} and ΔP_{aero}.

The relationship between inferred and actual total power difference can be seen in Table 4.12. To aid in understanding, all the slopes in the table are represented graphically in Figure 4.16. The percentages in Table 4.12 represent the combined influence of ΔP_{aero} and ΔP_{AD} on ΔP_{true}.

Table 4.12: The slopes of the regression for absolute power and fuel differences, $\Delta P_{true} \sim \Delta P_{inferred}$ and $\Delta F_{true} \sim \kappa \Delta P_{inferred}$. For fuel results, η_{gen} was set to 32.2%. Values in **bold** are those that are not significantly different from 100%, according to the 95% confidence interval.

Regression	$\Delta P_{true} \sim \beta_0 + \beta_1 \Delta P_{inferred}$			$\Delta F_{true} \sim \beta_0 + \kappa(\beta_1 \Delta P_{inferred})$		
$\bar{P}_{AD}/\bar{P}_{aero}$	None	RLS	C.Offset	None	RLS	C.Offset
Schmid	**98.9** $\pm$ 1.4	107.5 $\pm$ 1.6	104.4 $\pm$ 1.5	98.4 $\pm$ 1.3	**100.8** $\pm$ 1.5	**99.1** $\pm$ 1.5
Hussein RP	89.2 $\pm$ 1.3	93.7 $\pm$ 1.5	92.0 $\pm$ 1.4	88.3 $\pm$ 1.3	86.5 $\pm$ 1.5	85.9 $\pm$ 1.4
Hussein Pwr	91.1 $\pm$ 1.4	96.9 $\pm$ 1.6	94.9 $\pm$ 1.5	90.4 $\pm$ 1.4	90.8 $\pm$ 1.6	89.9 $\pm$ 1.5

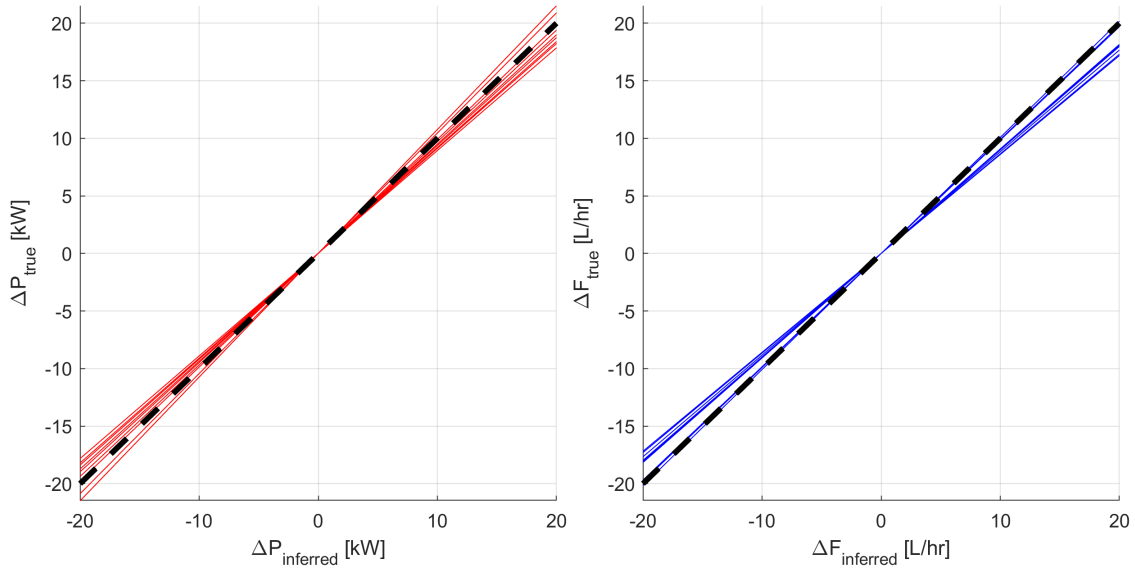

Figure 4.16: The slopes of the various methods in absolute terms ΔP and ΔF. The one-to-one line is also shown as a dotted black line.

In the combination of braking loss and drag reduction to create a net $\Delta P_{inferred}$, some clarity is lost: Certain methods that were recognized as biased with separate aero and

braking terms, now produce slopes that are indistinguishable from 100% when combined. This occurs because if the estimates for $\bar{P}_{AD}$ are biased in the opposite direction to the estimates for $\bar{P}_{aero}$, the two will average out. For example, the unadjusted $\bar{P}_{AD}$, originally around 90%, and the Schmid drag reduction model, with a coefficient of approximately 110%, when combined, yield a slope of roughly 100%. Consequently, a misleading impression of a good fit is generated when, in fact, neither braking loss nor drag is faithfully represented.

As for the relative outcomes ($NPC_{inferred}$ and $NFC_{inferred}$), the narrative closely resembles the absolute results. There is only one single coefficient of note in the relative results, that is, the ratio of inferred benefits to actual observed benefits. Table 4.13 shows this coefficient, in the now familiar form of a percentage, and Figure 4.17 displays the $NPC_{inferred}$ and $NFC_{inferred}$ slope coefficients graphically.

Table 4.13: The slope of each method on $NPC_{true} \sim NPC_{inferred}$ and $NFC_{true} \sim NFC_{inferred}$ (expressed as a percentage). For fuel results, η_{gen} was set to 32.2%. Values in **bold** are those that are not significantly different from 100%, according to the 95% confidence interval.

Regression	$NPC_{true} \sim NPC_{inferred}$			$NFC_{true} \sim NFC_{inferred}$		
$\bar{P}_{AD}/\bar{P}_{aero}$	None	RLS	C.Offset	None	RLS	C.Offset
Schmid	96.7 ± 1.5	105.1 ± 1.6	101.8 ± 1.6	94.9 ± 1.4	$\mathbf{98.8} \pm 1.7$	96.3 ± 1.6
Hussein RP	88.4 ± 1.4	92.9 ± 1.5	91.2 ± 1.5	85.9 ± 1.4	86.6 ± 1.6	85.6 ± 1.5
Hussein Pwr	89.8 ± 1.5	95.7 ± 1.6	93.6 ± 1.6	87.1 ± 1.4	89.1 ± 1.7	87.8 ± 1.6

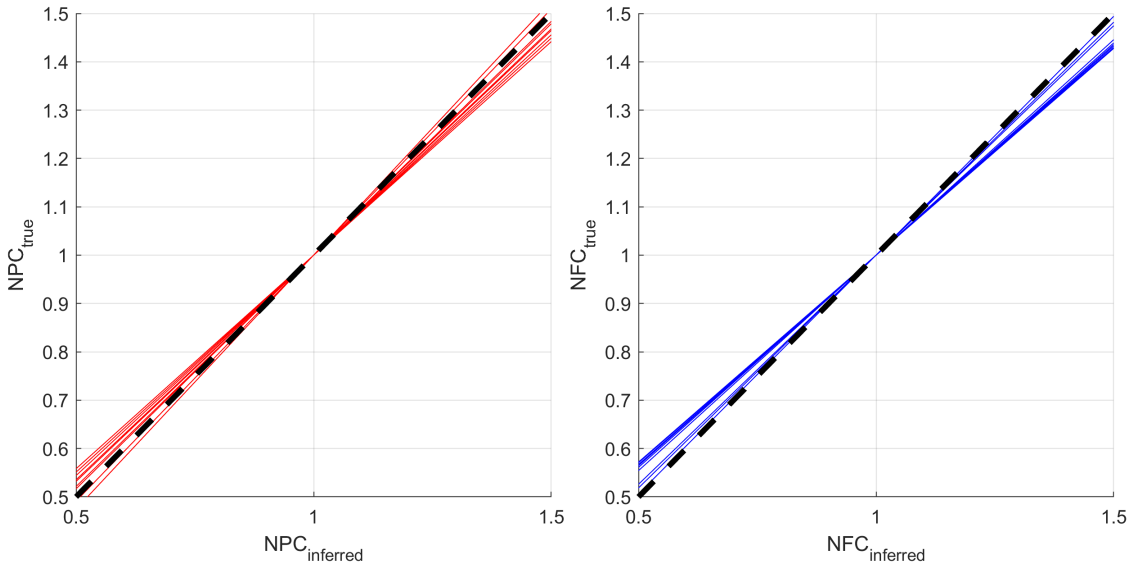

Figure 4.17: The slopes of the various methods in relative terms NPC and NFC. The one-to-one line is also shown as a dotted black line.

For power inferences, not one of the percentages coincides with 100%, although the Schmid model with Constant Offset comes very close. The RLS-Schmid fuel result overlaps 100% at the chosen κ. Most of the results are overpredictive of $NPC_{inferred}$, leading to coefficients less than 100%. This can be attributed in part to the optimistic prediction of drag savings of the Hussein models.

This concludes the presentation of the experimental results. A general discussion follows to discuss the merits and limitations of the proposed framework for predicting platoon energy benefits. After this discussion, the chapter will conclude with a rudimentary platoon savings classifier to demonstrate a potential use case for this framework.

4.2.6 Discussion

What is the significance of the results that have just been presented? First, it is important to acknowledge the reality of the experimental data. Energy inferences have been presented in comparison to a supposed "true" value. This depiction is not entirely accurate. The power and fuel signals that underlie the presumed differences are not actual measurements, but rather computed signals transmitted through the vehicle CAN network. Consequently, the purported "true" disparities in power and fuel contain inaccuracies. Moreover, it is improbable that these inaccuracies follow a normal distribution or are centered around zero. Nevertheless, the "true" values offer insights into the energy consumption of the vehicles, originating from a distinct source compared to the inferences, and are worth examining.

The value of κ is extremely important for the fuel results. It does not affect the correlation coefficients at all, but it directly impacts the slope. It is a bold assumption that a single characteristic value of κ can be effective; the trucks in this experimental dataset have different powertrain designs, with three different engine designs between the four of them. A κ value of $0.322 \cdot 36\text{e}6$ $^{\text{L}}/_{\text{J}}$ seemed reasonably appropriate, but a higher-fidelity powertrain model would improve things here.

There is also the issue of data imbalance. 82% the comparisons come from Dataset 1, since that is the most comprehensive set available. As a result, the regression results are most informed by the information in that dataset, which means heterogeneous platoons of two and four trucks, with unladen trailers, running at 45 mph (20.1 m/s). However, the agreement with other data sets was substantial and the unloaded trailers led to a highway-relevant drag fraction.

Next, there is the effect of the pairwise model matrix. There are two clear issues introduced by this matrix design. First, the matrix introduces a mixed effect (specifically, a multiple-membership effect) that has not been accounted for. Second, it yields slightly different results if the groups are numbered differently since the combinations are unordered,

but the elements of the model matrix are ordered. Neither of these drawbacks was severe enough to preclude the use of pairwise design. Modeling the mixed effects would serve to increase the precision of the slope estimates and may provide more insight about the adequacy of the various submodels for drag reduction and braking. This is left to future work.

For various reasons, it would be misguided to apply the fitted experimental regression slopes as a kind of "correction factor". This would bring empiricism into a physics-based method, in addition to the challenge of distinguishing between signal noise and modeling error due to the uncertainty of the "true" signals. The focus should be on creating accurate drag reduction models and braking power estimates, not fixing the biases of models post-hoc based on linear regression results. The models tested here may be good enough to apply in practice: the Schmid drag reduction model is somewhat conservative but agrees well, and the evidence indicates that the Constant Offset routine does an adequate job of adjusting $\bar{P}_{AD}$. Even if the models here are not "good" enough (in whatever sense of the word), the framework is still a demonstrated success; it may simply need better models.

We have demonstrated that the approach presented in this thesis produces predictions that exhibit a strong association with real variations in energy usage in various scenarios. These scenarios include:

- various types and models of trucks,
- both on public roads and controlled test tracks,
- with and without other vehicles nearby,
- under different weather conditions, and
- with and without instances of vehicles merging into the platoon.

With the robust regression technique that was used, up to 70% of the variance in platoon energy change was explained. The impact of accessory loads such as engine fans and HVAC subsystems was also not included here, due to the lack of information regarding these loads.

There is good reason to believe that with the appropriate mixed effects modeling strategy and accessory load modeling, the explained variance would be significantly higher.

Overall, the following assertions are made:

- the proposed methods for braking losses and drag reduction show good agreement with the measured energy changes, both on a fuel and power basis. The best of the methods may be within 2% or less on a relative basis.
- the Schmid drag model yielded significantly lower drag reduction estimates than the Hussein drag models, making it the most conservative of the drag reduction models
- the RLS and Constant Offset routines yielded lower braking estimates than unadjusted $\bar{P}_{AD}$. Of the two, the Constant Offset method appears superior. Both adjusted estimates were usually closer to 100% than the unadjusted estimate.
- the approximate technique for modeling the effect of wind from queried weather data did not add significant explanatory power
- the model could be used to predict platooning energy consumption during a microtrip in real time with reasonable confidence, given the R^2 of approximately 0.7.

Next, a potential and important use case of the model will be shown: the go/no-go decision of whether platooning saved energy. This is one of the most feasible applications of this dissertation and indeed the minimum criterion for the framework to be deemed useful. If a classifier can be created that is correct more than 50% of the time, then useful information is provided to the operator. More than that, we will propose a three-way classification scheme that incorporates a "maybe" category, using the estimated probability of the classifier.

4.3 A Rudimentary Platoon Energy Savings Classifier

Based on the results, it appears that we could use the framework to infer whether or not a platoon has saved fuel. The simplest scheme possible for making this "go/no-go" decision would be to assume that if the inferred energy change is less than zero (or the ratio less than one), then the platoon has saved energy. This simple scheme only incorporates binary information and uses none of the information about how close to the boundary a datapoint is, so it is hardly the best we can do. However, it provides an understandable entry point into the topic.

One common way to represent the output of classification algorithms is to place the count of predicted values versus their true values into a "confusion matrix". A confusion matrix displays the count of correct classifications along the left diagonal and the count of false positives and negatives along the right diagonal. Using our very simple classification scheme, Figure 4.18 displays the confusion matrix using $NFC > 1$ as the criterion. There was no "training" that occurred to create the outcomes in the confusion matrix in Figure 4.18, since it remains a physics-based guess at the energy change between two comparable platooning results. All datasets except Dataset 2 were used to construct Figure 4.18. Dataset 2 was excluded in the interest of maintaining a conservative estimate of accuracy, because the vast majority of that dataset contains no braking and it could lead to a false sense of high accuracy. Using Datasets 1, 3, and 4, 74.3% of the platooning runs were accurately classified.

74.3% is much better than randomly guessing, but not impressively accurate. However, the "true" class itself includes many errors, and there is a concentration of points about the classification threshold ($NFC = 1$). As such, 74.3% accuracy is acceptable. Interestingly, the better classification accuracy is obtained by using the headwind/air density estimates, which did not improve the regression correlation earlier in the chapter. Without including headwind, the accuracy is only 70%.

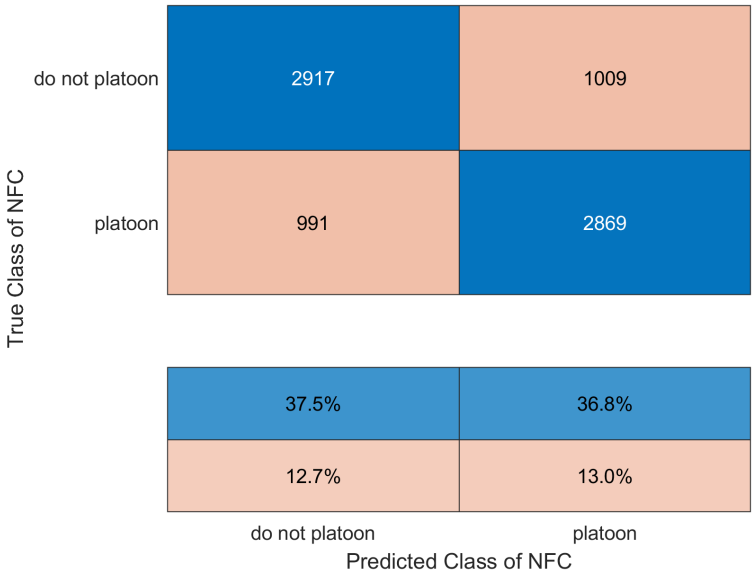

Figure 4.18: The confusion matrix for binary platooning classification using $NFC_{inferred} > 1$ as the predictor.

It is informative to see which trucks are easiest to predict. Looking at the truck-by-truck results in Figure 4.19, we see that the truck A1 was the hardest to predict, with an accuracy of 69.2%, and that the other trucks were predicted with greater than 75% accuracy. The reason why the energy change of A1 is classified less accurately is no doubt because A1 only experienced modest aerodynamic benefits since it always led the platoons it participated in. Similarly, military trucks (T13 and T14) were easier to classify due to their poor aerodynamics and higher mass, both of which increase their sensitivity to platooning energy change.

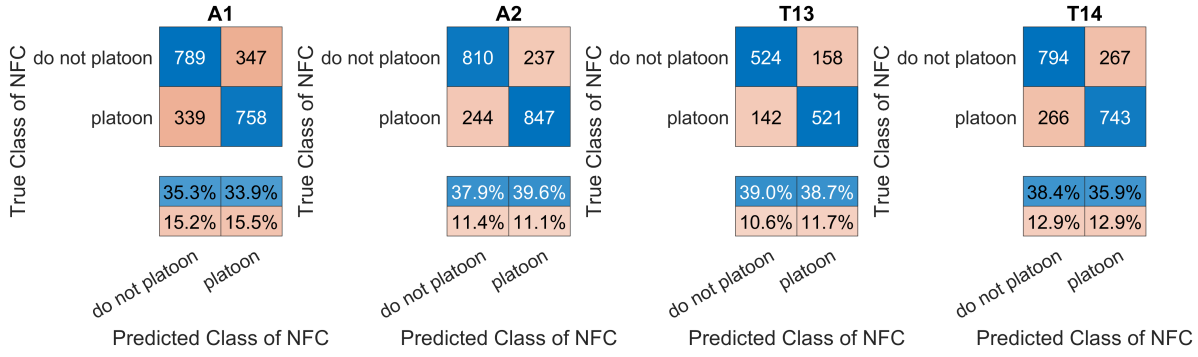

Figure 4.19: The confusion matrices broken out by truck for binary platooning classification using $NFC_{inferred} > 1$ as the predictor.

4.3.1 Logistic Regression to Provide Estimated Probability of Savings

To build a better classifier of platooning energy change, one that incorporates probability, we will turn to logistic regression. Logistic regression is the linear regression of a logit function, with the output representing an estimated probability of a given class[6]. The form of our logistic regression will use one or more inferred savings measures ($NFC_{inferred}$, $\Delta P_{inferred}$, etc.) to predict the probability of actual savings, as shown generally by Equation 4.6.

$$\Pr(\text{actual savings}) = \frac{1}{1 + e^{-(\beta_0 + \sum \beta_i (\text{inferred savings}))}} \tag{4.6}$$

The form of the logistic regression function is quite useful for mapping continuous inputs to a probability of an outcome since it is bounded from zero to one and smoothly varies between the two. Logistic regression coefficients are fit numerically, often using Maximum Likelihood Estimation [118].

We have several options for specifying the model formula of the logistic regression. A few candidate models to insert in Equation 4.6 are listed below in Table 4.14. The simplest logistic regression we can build is one where $NPC_{inferred}$ or $\Delta\bar{P}_{inferred}$ is the sole predictor of whether differential platoon energy consumption was positive or negative[7]. We can increase

Table 4.14: Six potential logistic regression models for platoon savings classification.

LR Model ID	Regressors	Formula
A	$NPC_{inferred}$	$\Pr(NFC_{true} < 1) = 1/1 + e^{-(\beta_0 + \beta_1 NPC_{inferred})}$
Af	$NPC_{inferred}$, $\Delta\bar{P}_{fan}$	$\Pr(NFC_{true} < 1) = 1/1 + e^{-(\beta_0 + \beta_1 NPC_{inferred} + \beta_2 \Delta\bar{P}_{fan})}$
B	$\Delta\bar{P}_{inferred}$	$\Pr(NFC_{true} < 1) = 1/1 + e^{-(\beta_0 + \beta_1 \Delta\bar{P}_{inferred})}$
Bf	$\Delta\bar{P}_{inferred}$, $\Delta\bar{P}_{fan}$	$\Pr(NFC_{true} < 1) = 1/1 + e^{-(\beta_0 + \beta_1 \bar{P}_{inferred} + \beta_2 \Delta\bar{P}_{fan})}$
C	$\Delta\bar{P}_{aero}$, $\Delta\bar{P}_{AD}$	$\Pr(NFC_{true} < 1) = 1/1 + e^{-(\beta_0 + \beta_1 \Delta\bar{P}_{aero} + \beta_2 \Delta\bar{P}_{AD})}$
Cf	$\Delta\bar{P}_{aero}$, $\Delta\bar{P}_{AD}$, $\Delta\bar{P}_{fan}$	$\Pr(NFC_{true} < 1) = 1/1 + e^{-(\beta_0 + \beta_1 \bar{P}_{aero} + \beta_2 \bar{P}_{AD} + \beta_3 \Delta\bar{P}_{fan})}$

the model flexibility of the classifier to adapt to class separation by individually using the

[6]For further reading on logistic regression and other classifiers, see [117]

[7]We could also use the equivalent fuel terms ($NFC_{inferred}$ or $\Delta\bar{F}_{inferred}$), though the result would be very nearly the same under the assumptions made in this work.

power differences as the predictors. We also have an estimated fan power difference (see Appendix A), which may add further explanatory power.

The models which use $NPC_{inferred}$ (A, Af) will more equally weight the differences in consumption between trucks than the models using power deltas (B, C, Bf, Cf). For the models with fan power (Af, Bf, Cf), the fan power term will be used to drive down the variance in the true classes where fan power differed among the selected pair. The situations where the power draw of the engine fan can be attributable to platooning are presumably rare (though such concerns have been raised in the literature), so in application, the expected contribution of the fan power term would be zero. The fan power term serves to explain some of the variance in the "true" class, but it would rarely contribute to a real-time inference.

Model Selection

Before selecting a logistic regression model structure from the options in Table 4.14, we first divide the data into training and test sets. There are various ways in which we could accomplish this division. The training sets will be Dataset 1, as it includes all four trucks and has good excitation in both braking and drag reduction. The test set will consist of the on-road datasets, Dataset 3 and 4. Dataset 2 will not be used for reasons that were previously discussed.

To quickly isolate the most promising logistic regression model, we will train the classifier for each of the 18 methods previously described, and look at the average test accuracy of each. For instance, for model A, we will train a logistic regression classifier to classify the training data based on $NPC_{inferred}$ 18 times, once for each combination of the weather, $\bar{P}_{AD}$ adjustment, and DRR methods. The average test accuracy among those 18 logistic regression models will be recorded for logistic regression model A, and compared to the average accuracy for model B, C and so on.

Table 4.15 lists the resulting average and best test accuracy for each logistic regression model ID. From an inspection of Table 4.15, we can see that the inclusion of estimated power

increased the training accuracy by a few percent, but had no impact on the testing accuracy.

Table 4.15: Accuracy of the logistic regression model candidates, averaged across the 18 combinations of DRR, $\bar{P}_{AD}$ adjustment, and weather models. Using probability threshold of 0.5.

LR Model ID	Training Accuracy		Testing Accuracy	
	Average	Max	Average	Max
A	72.43%	74.89%	71.32%	74.17%
Af	75.63%	78.72%	**71.47%**	74.17%
B	72.41%	74.93%	71.31%	**74.28%**
Bf	75.69%	78.70%	71.41%	**74.28%**
C	72.62%	75.62%	71.03%	73.31%
Cf	**75.89%**	**79.25%**	71.14%	73.31%

In terms of test accuracy, the best performing logistic regression model is model B or Bf, although models A and Af shows very similar performance. With model B and Bf, a test accuracy of 74.28% is achieved using the Schmid DRR model, no $\bar{P}_{AD}$ adjustment, and with the rough approximation for queried weather data differences. Interestingly, the modeling of air density and headwinds improved the test accuracy by 3-4% for all methods. The unadjusted $\bar{P}_{AD}$ estimate should probably be avoided, given its sensitivity to parameter errors. Substituting the Constant Offset adjustment routine does slightly decrease test accuracy to 73.42%, but this is an acceptable tradeoff for more robust estimation. The chosen logistic regression model that will be analyzed, then, is model B, which is a one-term model with $\Delta\bar{P}_{inferred}$ as the predictor. We will use the Schmid DRR model, Constant Offset $\bar{P}_{AD}$ adjustment, and queried weather corrections to calculate $\Delta\bar{P}_{inferred}$. However, there is an argument to be made that using NPC_{inf} (model A or Af) instead would be more generalizable, as it is a relative measure that allows cross-comparison of different truck's benefit. Therefore, we will also check model A's performance to see if it yields better test accuracy.

The normalized model B logistic regression values for slope and intercept[8] for $\Pr(NFC_{true})$ versus $\Delta\bar{P}_{inferred}$ are -2.2764 (*p-value: 3.6e-261*) and -0.0351 (*p-value: 0.2287*) when using Schmid, Constant Offset, and headwind. The intercept (β_0) is not significant at a 95% confidence level, meaning that there is little evidence that the true savings are divided along a power difference other than zero. If the classes were less separated, then β_1 would have a lower value. For instance, the normalized slope coefficient for model A is lower: -2.1975 (*p-value: 2.8e-267*). This indicates that there is more class overlap when $NPC_{inferred}$ is used instead of $\Delta\bar{P}_{inferred}$.

Adding the Outcome "Maybe"

If pure two-way classification is desired, then the go/no-go decision would presumably be made on the 50% probability line. To incorporate the possibility of "maybe", we can specify a loss rate (say, 5%) and find an upper and lower probability that leads to that loss rate, placing everything in between as a "maybe". For smaller and smaller loss rates, the proportion of points that are categorized as "maybe" will increase until everything is a maybe. This loss rate applies to the training set, but not necessarily the testing set, which may have different levels of class overlap than the training set.

A common method for visualizing the loss rate is the receiver operating characteristic (ROC) curve, which shows a classifier's true positive rate versus its false positive rate, with the worst possible classifier being a one-to-one line. Figure 4.20 shows the ROC curve for a logistic regression classifier applied to both the training set and the test set. The area under the curve (AUC) is a measure of classification performance, with higher AUC being better, up to a max of 1. The AUC is 0.8307 for the training set and 0.8318 for the test set, which indicates that the test set is being classified at a similar accuracy as the training set. One way to understand AUC is as the probability that for a pair of true positive and negative

[8]The values for slope/intercept are stated after centering and scaling $\Delta\bar{P}_{inferred}$ by 6 W and 2288.1 W respectively. This process, known as normalizing, transforms the target variables to zero mean and variance of one. While it was not absolutely necessary here, it is best practice for logistic regression tasks.

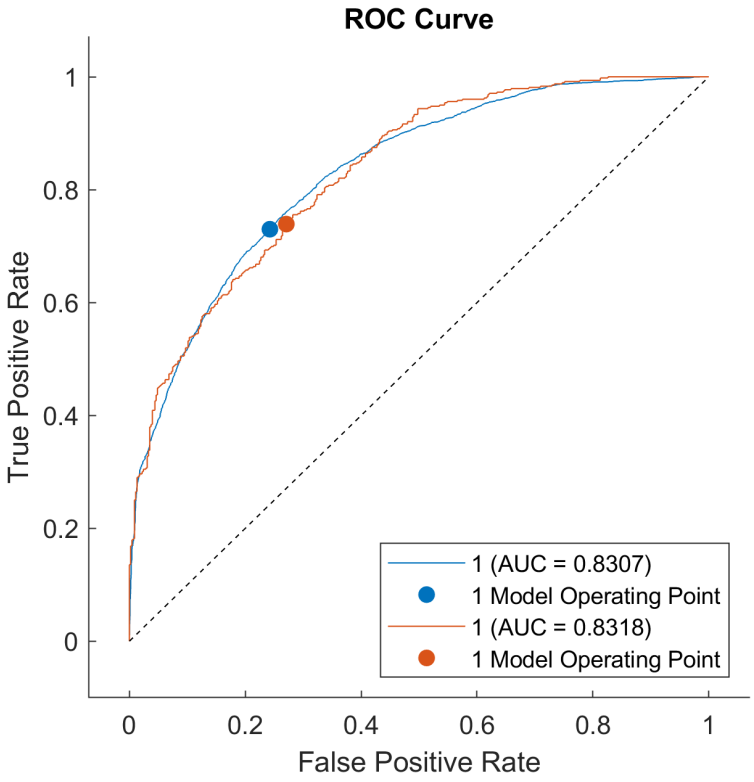

Figure 4.20: The ROC curve for the logistic regression model structure B ($\Pr(NFC_{true} < 1)$ versus $\Delta\bar{P}_{inferred}$) for the training and test set. The true class is equivalent to $NFC_{true} < 1$ in this case.

results, the output for the positive result would be higher than that of the negative result [117]. Taken this way, about 83% of the time, a true positive fuel saving would have a higher estimated probability than a true negative fuel saving.

Using the same information that was used to construct the ROC curve, we can arrive at the thresholds for where the lines for yes/no/maybe should be drawn. We first specify an acceptable loss rate (5%), then find the probability that corresponds to that loss rate for both the "saved fuel" class and the "did not save fuel" class. For example, with an allowed loss rate of 10%, the lower and upper probabilities are 37% and 62.5%.

The confusion matrix resulting from this 10% loss rate is shown in Figure 4.21, where it can be seen that 10% of the training samples are misclassified, 52.4% are correctly classified, and the remaining 37.6% are deemed "Maybe". In the test set, the accuracies are different, with 20.2% total misclassification, and 14.3% "Maybe". This emphasizes that the specified

loss rate is with respect to the training set. The lack of the class "Maybe" in the true classes is because we have no way to tell if the true class is a maybe. Interestingly, using model A (based on $NPC_{inferred}$) instead of model B, the total test misclassification is 16.4%, which is closer to the targeted 10%. This hints that $NPC_{inferred}$ may be a better variable for classification than the pure inferred power difference $\Delta\bar{P}_{inferred}$.

As a final point, Figure 4.22 shows the decision boundaries for platooning based on the trained logistic regression, this time using model structure A. It can be seen that the width of the class "Maybe" decreases as the desired level of confidence decreases, until it converges at the two-way decision bound. The bounds are believed to be relatively conservative, due to the unaccounted for variance from the pairwise model matrix and the "true" class.

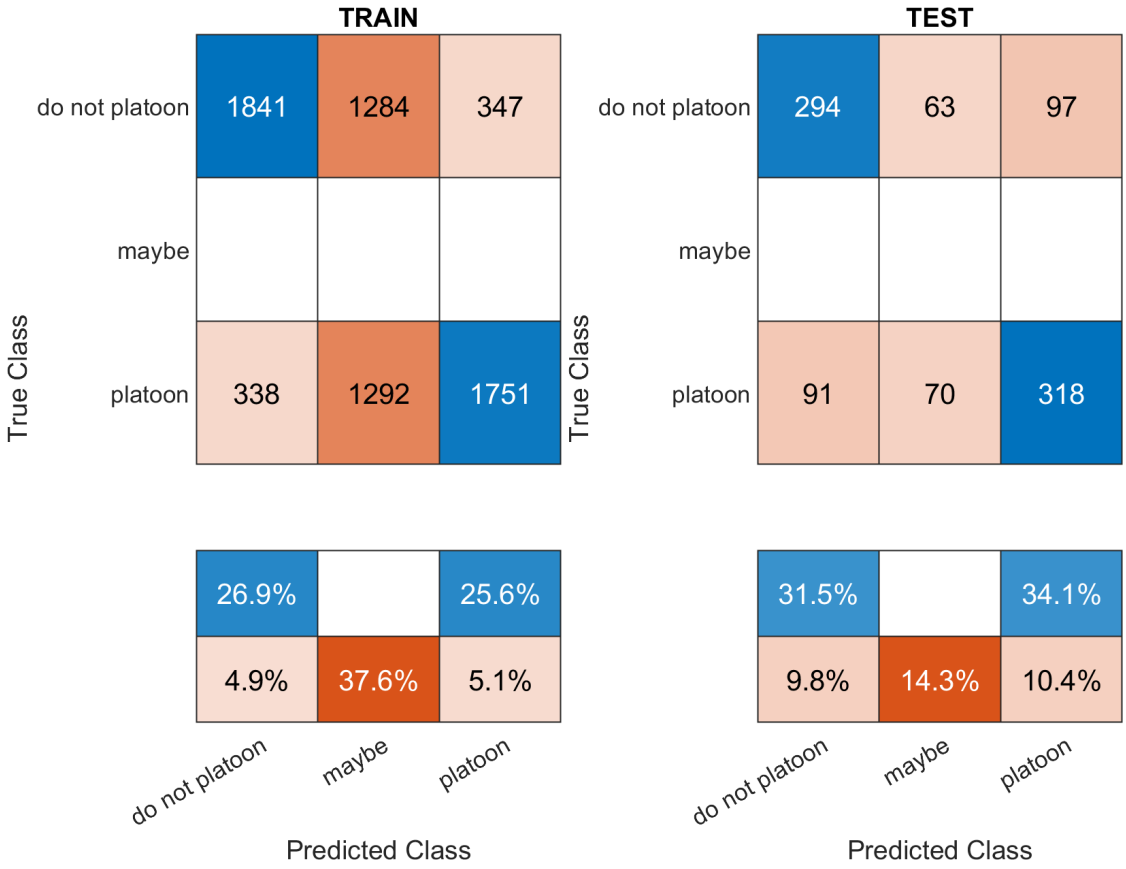

Figure 4.21: The three-way confusion matrix for the logistic regression model structure B ($\Pr(NFC_{true} < 1)$ versus $\Delta\bar{P}_{inferred}$) for both the training and test set, with specified acceptable loss rate of 10%.

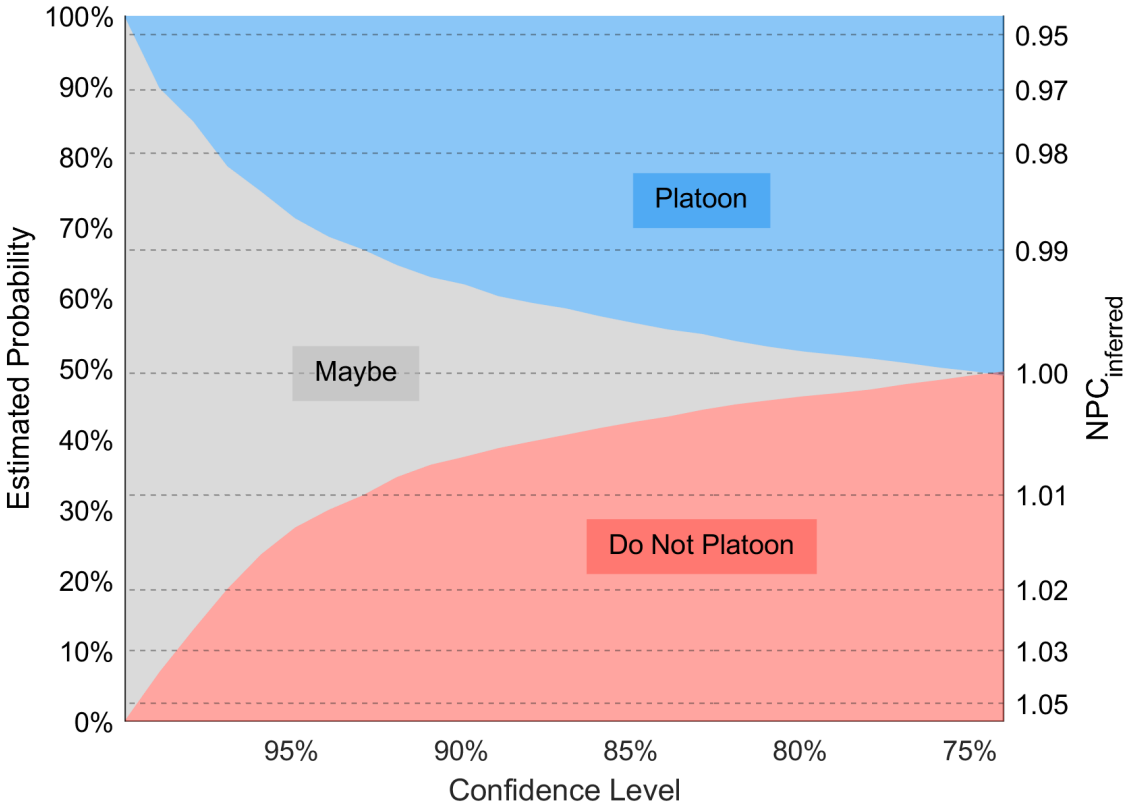

Figure 4.22: Trained platoon decision by logistic regression, based on desired two-sided confidence level. The estimated probability of fuel savings is mapped to the left axis and the levels of $NPC_{inferred}$ that correspond to the probabilities are mapped to the right axis.

Figure 4.22 is an extremely practical result. Given inferred fuel consumption (which can be readily calculated in real time), it can tell an operator whether they benefitted from platooning at a chosen level of confidence. For example, fuel savings were observed for about 95% of the DOE data with an $NPC_{inferred}$ below 0.985 (1.5% benefit).

Overall, this is merely an application of the framework in this dissertation to draw some conclusions about whether a platoon saved energy. The classification methods here are preliminary, but even still, the schemes do a passable job of making the go/no-go decision, with a 73.5% accuracy for the on-road platooning data in Datasets 3 and 4.

It cannot be overemphasized that the classifier proposed here is trained entirely on the DOE data, and the actual observed "go/no-go" response itself contains plenty of uncertainty. The observed value for "go/no-go" can hardly be called a true value, due to the pairwise

comparison method, lack of accounting for weather and accessory load differences, inaccuracy of the engine fuel rate signal, and so on. In practice, the classification task would be easier than here: Instead of classifying the difference between two separate instances of a truck's fuel consumption, only a single instance would be classified, which is less susceptible to error than a comparison.

In summary, several methods were explored in this section to classify platoon performance. A logistic regression classifier was trained on the DOE dataset (Dataset 1). This classifier was then used to demonstrate a three-way classification technique, by which operators can specify their desired level of confidence in whether energy savings were achieved. The classifiers discussed in this section serve as a solid foundation for employing the framework presented in this dissertation to categorize platoon performance.

4.4 Chapter Summary

The chapter began with a sensitivity analysis, which applied errors to wheelspeed, grade, and engine power, and induced uniform errors in vehicle mass, drag coefficient, and rolling resistance coefficient. Overall, the platoon energy prediction framework was relatively robust to drag and rolling resistance errors, but mass errors did skew the results by a two-to-one ratio when using Constant Offset or RLS to estimate $\bar{P}_{AD}$.

The case study of the Canadian cut-in data was then revisited to see if the proposed framework worked better than previous methods (Section 3.2). In many cases, the predicted $NPC_{inferred}$ fell within the published experimental error bounds, indicating that the framework could also work for experimental data and that it successfully generalizes. The Constant Offset $\bar{P}_{AD}$ adjustment and the Schmid model were identified as good choices among the various submethods.

After that, a way to combine all four datasets into a single analysis was derived, which produced the model matrix. The pairwise design of the model matrix was seen to have various

strengths and weaknesses, which were discussed and partially allayed by the application of robust regression techniques.

Then, a complete representative result was generated, using the Constant Offset routine for $\bar{P}_{AD}$, the Schmid drag reduction model, and assuming constant weather conditions. The entire pairwise model matrix was used. Power differences (ΔP_{true}) were well captured, as were the relative differences (NPC_{true}). Fuel was more difficult to interpret, as the power-to-fuel conversion factor (κ) directly determines the relationship. Otherwise, the fuel results were very similar in terms of goodness-of-fit to the power results, although there was some indication that the constant powertrain efficiency assumption was being violated when the braking losses were significant.

After representative results were given, the results were expanded to all 18 possible methods for platoon energy inference. Ultimately, most of the methods demonstrated good agreement with the observed power and fuel differences. The recommended combination of methods is the one that was used for the representative result, which uses the Schmid drag reduction model, the Constant Offset routine for braking power correction, and assumes constant weather.

A discussion of the experimental findings followed. With the vast quantity of data and methods available, it was emphasized to avoid drawing conclusions beyond what the results actually show. In particular, it was recommended to avoid empirical adjustment of platooning inferences. Instead, it was proposed to focus on enhancing the submethods (such as the DRR models). Some simple enhancements to the submethods include considering accessory load modeling (as engine fans can consume significant power) and using real measured wind data instead of the estimated values that were utilized in this study.

The chapter ended with an exploration of platoon energy classification. Such classifications could be used to provide useful operator feedback, which would gamify platooning. A simple Boolean classification scheme was about 73% accurate, although concerns about the pairwise model matrix hold here; much of the misclassification could just be due to noise

in the measurements. The prediction intervals of the regression results were considered to be too conservative, which motivated the development of a dedicated classifier using logistic regression. The logistic regression model was trained on Dataset 1, and tested on the on-road datasets. A third class "Maybe" was incorporated into the classifier, which allows the user to specify their desired confidence level. Despite its demonstrative nature, the classifier as presented is a promising application of the framework.

Chapter 5

Conclusions

In recent years, platooning has faced skepticism due to uncertain energy benefits. This dissertation sought to counteract this trend and demystify the energy consumption of platoons. A robust physics-based framework was developed to infer platoon benefits in real time. In simulation and real data, it was seen that the inferences resulting from this framework were well correlated with the actual energy consumption of the platoons.

5.1 Contribution

The framework fulfills the requirements laid out in the Introduction:

- **Energy Impact Isolation:** It distinctly isolates energy changes attributable to platooning by using no-braking/no-drag-reduction conditions as the baseline. This precision is maintained even in complex scenarios, such as vehicle cut-ins and hilly terrain.
- **Adaptable:** It is modular and can be easily applied to future vehicle architectures, only requiring braking power, drag reduction power, and overall energy consumption, which are fundamental aspects of road vehicles.
- **Real-Time:** Apart from the slight delay caused by estimating acceleration, the techniques described are executed in real time and do not require a baseline reference.
- **Known Sensitivity:** It has a known sensitivity to errors, with notable resistance to rolling resistance errors, and a defined response to variations in vehicle mass and aerodynamic drag.
- **Feedback Mechanism:** It incorporates a feedback mechanism for binary responses, showing an initial test accuracy of 73% (yes/no) for complex and noisy real-world data sets. Additionally, it offers the option of setting a target accuracy level by introducing a "maybe" response, thereby increasing its usefulness in practical applications.

This dissertation introduces an innovative framework that pioneers the prediction of platooning energy effects with control-induced dynamics, a capability previously unaddressed

in the field. Furthermore, the practical applicability of this framework has been demonstrated on experimental data from many sources, using a novel algorithm to estimate vehicle braking power from acceleration.

Great emphasis has been placed on the practical implementation of the framework in real vehicles. The signals needed to properly implement the methods in this framework are as follows:

- engine power or fuel rate
- longitudinal velocity (from which acceleration can be estimated)
- a Boolean indication of active deceleration
- road grade
- V2V range

The required vehicle parameters include:

- weight
- transmission ratios and rotational moment of inertia (for effective mass)
- aerodynamic drag coefficient and frontal area, or drag-area value
- rolling resistance coefficient

To strengthen the practical credibility of the framework, it was applied to data from many different platoon trials and was found to be well-conceived and tractable. Secondary contributions of this thesis are the testing of several drag reduction models in the literature on data unseen to the models, as well as corrections to a published model's coefficients.

The framework presented in this thesis was conceived as a feedback mechanism for platoon operators. It is a method to provide real-time feedback to platoon operators about how much energy they have saved during platooning. By providing real-time inferences that are consistent with actual fuel savings, drivers are provided with a mechanism to concretely reward them for platooning.

5.2 Significance

A direct application of the framework in this dissertation would provide feedback to drivers and platoon coordinators about their differential platoon fuel savings. Ideally, the current platooning microtrip savings could be displayed alongside long-term cumulative energy savings from platooning. To illustrate this point, Figure 5.1 shows a head-up display (HUD) that could be displayed to a platoon operator using this framework. Both the savings of the present microtrip and the cumulative energy saved by platooning is shown in the HUD.

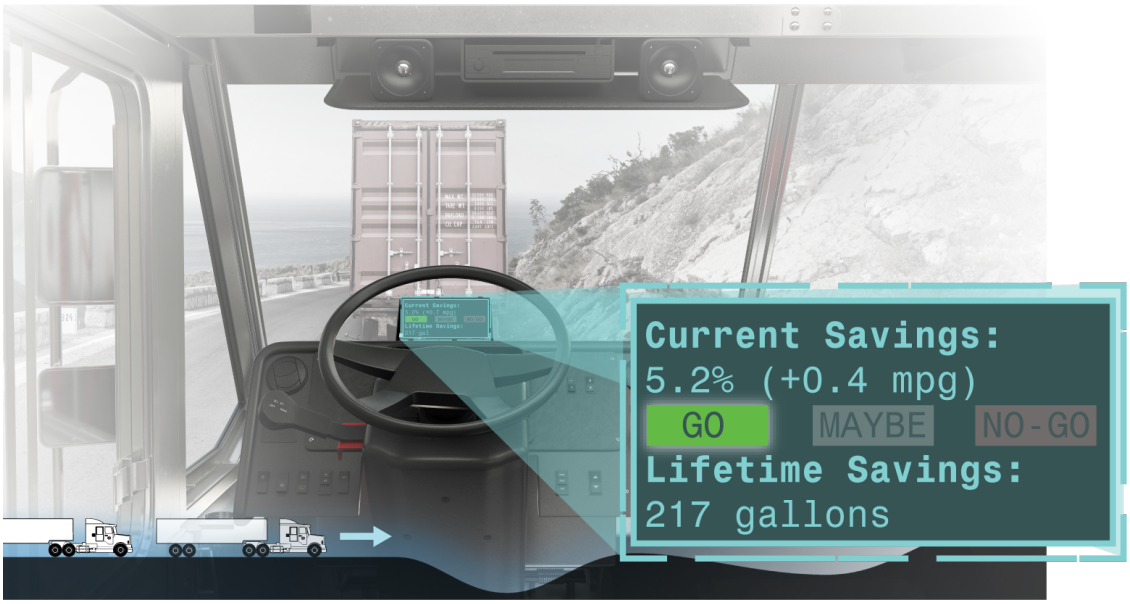

Figure 5.1: A demonstration of the kind of information that could be displayed to platoon operators by using this framework.

To show how the framework could be used to generate the information on the HUD in Figure 5.1, we can use an example. Suppose that the inferred benefit was 5.2% ($NPC_{inferred} = 0.948$), and the average fuel economy of the current platoon microtrip was 7.7 mpg. The inferred non-platooning mpg would be $7.7 \times 0.948 = 7.3$ mpg, which represents a 0.4 mpg increase in fuel economy. The gallons of diesel saved from there would be straightforward: One would take the number of miles traveled while platooning (say, a full shift

of 500 miles) and use the average mpg benefit to calculate the amount of fuel saved: $500 \times \left(\frac{1}{0.948 \times 7.7} - \frac{1}{7.7}\right) = 3.56$ gallons of fuel saved. It may take a long time to reach the displayed lifetime savings of 217 gallons, but this depends on the availability and motivation of drivers to participate.

For electric vehicles that take a long time to charge, the energy savings of platooning could be used to display the inferred cumulative charging time savings, providing drivers with a measure of the increased uptime they have generated by platooning. Take, for instance, an electric Freightliner eCascadia with a battery that has 350 kWhr of usable capacity and is capable of a 90-minute, 230 kW charging rate (specifications from [119]). If platooning was able to extend the eCascadia's range by five percent, from 220 to 231 miles, this would represent 5 minutes at the charger for a full duty cycle, in addition to the energy saved. While 5 minutes is seemingly minor, it is easy to see how if multiple trucks are lined up at a charger, these time savings could add up to a meaningful advantage.

Although human-driver feedback was the application for which it was conceived, the framework could also be used for optimization purposes. Unlike conventional optimization strategies that focus on end metrics, such as fuel consumption, this framework delves into the causal factors behind energy use. This approach provides insight into the "why" and "how" of platoon energy savings, paving the way for more effective energy management strategies.

While total energy optimization has its merits, particularly in its ability to encompass system nonlinearities and address complex driving scenarios, it often comes at the cost of interpretability. The proposed framework allows for a more granular attribution of energy consumption to specific operational conditions, especially active deceleration. Although this may lead to biased estimates of actual fuel costs, the diagnostic and interpretive benefits it offers are invaluable. Future research could explore methods to integrate the framework with the optimization of total energy consumption, which would provide a detailed picture of vehicle energy consumption from both the "cause" and "effect" ends.

5.3 Recommendations for Future Work

The framework presented herein marks a pioneering approach to estimating platoon energy savings, but it also opens avenues for refinement. Because it is a framework, the submodels for the calculation of drag reduction and braking loss can be easily replaced with improved submodels. In particular, the following recommendations for future work are made:

- **Improvements to the drag modeling:** The submodel for aerodynamic drag reduction is quite simple in this work. Not only was a rough approximation for headwind used, but the DRR models that were tested did not account for the effect of crosswinds on platoon drag reduction. A recent paper has introduced a way to account for the effects of crosswinds on effective DRR [95], which could be easily incorporated into the framework.
- **Reduced sensitivity to mass error:** The adaptive estimate of braking loss ($\bar{P}_{AD}$) was sensitive to errors in vehicle mass because of the grade-dependent acceleration error that mass causes. Future work could improve upon the robustness of the Constant Offset or RLS algorithms to decrease this mass sensitivity.
- **True real-time acceleration:** To get a linear-phase acceleration estimate for braking loss estimation, a time delay of 1.4 seconds was introduced into the signal processing chain. The inclusion of an acceleration measurement to achieve true real-time operation represents an immediate extension of this work that would make it more suited to control tasks.
- **Incorporation of braking effort values:** The acceleration-driven approach for $\bar{P}_{AD}$ estimation used here only uses a boolean braking status, which makes it vulnerable to errors during low-magnitude braking. To further improve the estimation of $\bar{P}_{AD}$, braking effort information could be leveraged, such as reported retarder torque or brake pressure.

- **Detailed powertrain modeling:** The simplification of powertrain modeling, such as constant conversion efficiency η_{gen}, highlights the need for more detailed powertrain modeling. An investigation into how platooning changes powertrain efficiency should be carried out.
- **Application to other vehicle architectures:** This work has been validated on HDVs with conventional turbodiesel powertrains and box trailers. Theoretically, this framework could apply to any platooning vehicle. The application of the framework to vehicles with regenerative braking systems represents a particularly promising area of study.
- **Application to leader braking:** While the majority of this work assumed a brakeless reference ($\bar{P}_{AD,\,ref} = 0$), the case of platooning where both leader and follower are braking was not explored in any depth. A way to handle braking losses not caused by platooning was introduced (see Equation 3.21b), but this should be examined in practice.
- **Higher quality truth data:** While the pairwise comparison process greatly increased the validation power of the experimental data, the process also inflated the variance of the experimentally derived values for true pairwise energy consumption differences (such as those in Section 4.2.4, Figure 4.12). The contribution of this variance is not known, and it would be interesting to see how the inference performs on data where the truth value for fuel economy is better known.
- **Field testing:** To date, the framework has been validated under nominally constant-speed operating conditions. The framework has not been implemented on a physical system, and its performance during significant speed variations remains to be explored. Field testing is a critical next step that offers the potential to illuminate the efficacy of the framework in diverse and unpredictable real-world scenarios.

5.4 Summary

In conclusion, this framework represents a significant step forward in the quest for viable platooning operations. The unsteady investment in platooning can be partly attributed to its unpredictable real-world savings. By providing realistic in-situ expectations of energy savings, this framework makes platoon energy savings tangible. It empowers operators with the knowledge to optimize platooning usage. Consequently, this model sets the foundation for a future in which platooning becomes an integral aspect of sustainable transportation.

"Let those who think I have said too little, or those who think I have said too much, forgive me; and let those who think I have said just enough join me in giving thanks to God."

- St. Augustine

References

[1] S. Davis and R. Boundy, "Transportation Energy Data Book (Edition 40)," Tech. Rep. ORNL/TM-2022/2376, 1878695, May 2022.

[2] U.S. Department of Transportation, Bureau of Transportation Statistics (BTS), Federal Highway Administration (FHWA), "Freight Analysis Framework Version 5," 2017.

[3] V. Milanes, S. E. Shladover, J. Spring, C. Nowakowski, H. Kawazoe, and M. Nakamura, "Cooperative Adaptive Cruise Control in Real Traffic Situations," *IEEE Transactions on Intelligent Transportation Systems*, vol. 15, pp. 296–305, Feb. 2014.

[4] A. Vahidi and A. Sciarretta, "Energy saving potentials of connected and automated vehicles," *Transportation Research Part C: Emerging Technologies*, vol. 95, pp. 822–843, Oct. 2018.

[5] V. A. Sujan, P. Jones, and A. Siekmann, "Characterizing Minimum Admissible Separation Distances in Heavy Duty Vehicle Platoons," in *2022 International Conference on Connected Vehicle and Expo (ICCVE)*, (Lakeland, FL, USA), pp. 1–6, IEEE, Mar. 2022.

[6] P. D. Brummitt and M. S. Khan, "Truck Platooning and Its Impact on Fuel Emissions," in *2022 IEEE Green Technologies Conference (GreenTech)*, (Houston, TX, USA), pp. 130–135, IEEE, Mar. 2022.

[7] A. K. Bhoopalam, N. Agatz, and R. Zuidwijk, "Planning of truck platoons: A literature review and directions for future research," *Transportation Research Part B: Methodological*, vol. 107, pp. 212–228, Jan. 2018.

[8] R. Stephens, "A review of gamified approaches to encouraging eco-driving," *Frontiers in Psychology*, vol. 13, p. 970851, Sept. 2022.

[9] J. Roberts, R. Mihelic, M. Roeth, and D. Rondini, "Confidence Report on Two-Truck Platooning," tech. rep., North American Council for Freight Efficiency (NACFE), Sept. 2016.

[10] E. Stegner, P. Snitzer, D. Bevly, and M. Hoffman, "New Metrics for Quantifying the Energy Efficiency of Platoons in the Presence of Disturbances," SAE Technical Paper 2022-01-0526, Mar. 2022.

[11] E. Stegner, P. Snitzer, J. Bentley, D. M. Bevly, and M. Hoffman, "Quantifying the Energy Impact of Autonomous Platooning-Imposed Longitudinal Dynamics," SAE Technical Paper 2023-01-0896, Apr. 2023. DOI:.

[12] J. Axelsson, "Safety in Vehicle Platooning: A Systematic Literature Review," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, pp. 1033–1045, May 2017.

[13] SAE J1321, "Fuel Consumption Test Procedure - Type II," Ground Vehicle Standard J1321_202010, SAE International, Oct. 2020.

[14] L. Zhang, F. Chen, X. Ma, and X. Pan, "Fuel Economy in Truck Platooning: A Literature Overview and Directions for Future Research," *Journal of Advanced Transportation*, vol. 2020, pp. 1–10, 2020.

[15] V. Turri, B. Besselink, and K. H. Johansson, "Cooperative Look-Ahead Control for Fuel-Efficient and Safe Heavy-Duty Vehicle Platooning," *IEEE Transactions on Control Systems Technology*, vol. 25, no. 1, pp. 12–28, 2017.

[16] J. Törnell, *Aerodynamics of vehicle platooning*. Dissertation, Chalmers University of Technology, Göteborg, Sweden, 2023.

[17] K. Salari and J. Ortega, "Experimental Investigation of the Aerodynamic Benefits of Truck Platooning," SAE Technical Paper 2018-01-0732, Apr. 2018.

[18] M. P. Lammert, K. J. Kelly, and J. Yanowitz, "Correlations of Platooning Track Test and Wind Tunnel Data," Tech. Rep. NREL/TP–5400-68897, 1422885, Feb. 2018.

[19] M. Zabat, N. Stabile, S. Frascaroli, and F. Browand, "The Aerodynamic Performance of Platoons: Final Report," Tech. Rep. UCB-ITS-PRR-95-35, University of Southern California, Los Angeles, CA, USA, Oct. 1995.

[20] J. Törnell, S. Sebben, and D. Söderblom, "Influence of Inter-Vehicle Distance on the Aerodynamics of a Two-Truck Platoon," *International Journal of Automotive Technology*, vol. 22, pp. 747–760, June 2021.

[21] M. P. Lammert, A. Duran, J. Diez, K. Burton, and A. Nicholson, "Effect of Platooning on Fuel Consumption of Class 8 Vehicles Over a Range of Speeds, Following Distances, and Mass," *SAE International Journal of Commercial Vehicles*, vol. 7, no. 2, pp. 626–639, 2014.

[22] SAE J1343, "Information Relating to Duty Cycles and Average Power Requirements of Truck and Bus Engine Accessories," Ground Vehicle Standard J1343_200008, SAE International, Aug. 2000.

[23] M. Ellis, J. I. Gargoloff, and R. Sengupta, "Aerodynamic Drag and Engine Cooling Effects on Class 8 Trucks in Platooning Configurations," *SAE International Journal of Commercial Vehicles*, vol. 8, pp. 732–739, Sept. 2015.

[24] S. N. P. Vegendla, T. Sofu, R. Saha, L.-K. Hwang, and M. Madurai Kumar, "Investigation on Underhood Thermal Analysis of Truck Platooning," *SAE International Journal of Commercial Vehicles*, vol. 11, pp. 5–16, Mar. 2018.

[25] B. Marcu and F. Browand, "Aerodynamic Forces Experienced by a 3-Vehicle Platoon in a Crosswind," SAE Technical Paper 1999-01-1324, Mar. 1999.

[26] B. McAuliffe, M. Lammert, X.-Y. Lu, S. Shladover, M.-D. Surcel, and A. Kailas, "Influences on Energy Savings of Heavy Trucks Using Cooperative Adaptive Cruise Control," SAE Technical Paper 2018-01-1181, Apr. 2018.

[27] K. Tadakuma, T. Doi, M. Shida, and K. Maeda, "Prediction formula of Aerodynamic Drag Reduction in Multiple-Vehicle Platooning Based on Wake Analysis and On-Road Experiments," *SAE International Journal of Passenger Cars - Mechanical Systems*, vol. 9, pp. 645–656, Apr. 2016.

[28] M. Schmid, D. Liu, B. Eksioglu, N. Huynh, and G. Comert, "Prediction Model for Energy Consumption in Heavy-Duty Vehicle Formations," in *IISE Annual Conference Proceedings*, pp. 1–6, 2020.

[29] E. Stegner, J. Ward, J. Siefert, M. Hoffman, and D. M. Bevly, "Experimental Fuel Consumption Results from a Heterogeneous Four-Truck Platoon," SAE Technical Paper 2021-01-0071, Apr. 2021.

[30] J. Törnell, S. Sebben, and P. Elofsson, "Experimental investigation of a two-truck platoon considering inter-vehicle distance, lateral offset and yaw," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 213, p. 104596, June 2021.

[31] M. P. Lammert, B. McAuliffe, P. Smith, A. Raeesi, M. Hoffman, and D. Bevly, "Impact of Lateral Alignment on the Energy Savings of a Truck Platoon," SAE Technical Paper 2020-01-0594, Apr. 2020.

[32] B. R. McAuliffe and M. Ahmadi-Baloutaki, "A Wind-Tunnel Investigation of the Influence of Separation Distance, Lateral Stagger, and Trailer Configuration on the Drag-Reduction Potential of a Two-Truck Platoon," *SAE International Journal of Commercial Vehicles*, vol. 11, pp. 125–150, June 2018.

[33] B. McAuliffe, P. Smith, A. Raeesi, M. Hoffman, and D. M. Bevly, "Track-Based Aerodynamic Testing of a Two-Truck Platoon," SAE Technical Paper 2021-01-0941, Apr. 2021.

[34] C. Zhang, M. Lammert, and B. McAuliffe, "Impact of Lateral Alignment for Cooling Airflow during Heavy-Truck Platooning," SAE Technical Paper 2021-01-0231, Apr. 2021.

[35] M. Song, F. Chen, and X. Ma, "Organization of autonomous truck platoon considering energy saving and pavement fatigue," *Transportation Research Part D: Transport and Environment*, vol. 90, p. 102667, Jan. 2021.

[36] H. Cheng, Y. Wang, D. Chong, C. Xia, L. Sun, J. Liu, K. Gao, R. Yang, and T. Jin, "Truck platooning reshapes greenhouse gas emissions of the integrated vehicle-road infrastructure system," *Nature Communications*, vol. 14, p. 4495, Aug. 2023.

[37] X. Hu, C. Peng, J. Bao, Y. Jiang, and S. Li, "Characteristics of naturally formed semitruck platoons on interstate highways and their implications for truck platooning

field deployment," *Canadian Journal of Civil Engineering*, vol. 51, pp. 345–355, Mar. 2024. Publisher: NRC Research Press.

[38] S. Tsugawa, S. Jeschke, and S. E. Shladover, "A Review of Truck Platooning Projects for Energy Savings," *IEEE Transactions on Intelligent Vehicles*, vol. 1, pp. 68–77, Mar. 2016.

[39] F. Browand and M. Hammache, "The Limits of Drag Behavior for Two Bluff Bodies in Tandem," SAE Technical Paper 2004-01-1145, Mar. 2004.

[40] J. Ortega and K. Salari, "Experimental Investigation of the Aerodynamic Benefits of Truck Platooning: Two- and Four-Vehicle Platoons," pp. 2021–01–0942, Apr. 2021.

[41] T. Gheyssens and G. Van Raemdonck, "Effect of the Frontal Edge Radius in a Platoon of Bluff Bodies," *SAE International Journal of Commercial Vehicles*, vol. 9, pp. 371–380, Sept. 2016.

[42] B. McAuliffe, A. Raeesi, M. Lammert, P. Smith, M. Hoffman, and D. Bevly, "Impact of Mixed Traffic on the Energy Savings of a Truck Platoon," *SAE International Jounal Advances and Current Practices in Mobility*, vol. 2, no. 3, pp. 1472–1496, 2020.

[43] K. R. Cooper, "The Wind Tunnel Testing of Heavy Trucks to Reduce Fuel Consumption," *SAE Transactions*, vol. 91, pp. 4118–4130, Feb. 1982.

[44] National Research Council, "Vehicle Technologies for Reducing Load-Specific Fuel Consumption," in *Technologies and Approaches to Reducing the Fuel Consumption of Medium- and Heavy-Duty Vehicles*, pp. 91–130, Washington, D.C.: National Academies Press, 2010.

[45] A. H. Taylor, M. J. Droege, G. M. Shaver, J. A. Sandoval, S. Erlien, and J. Kuszmaul, "Capturing the Impact of Speed, Grade, and Traffic on Class 8 Truck Platooning," *IEEE Transactions on Vehicular Technology*, vol. 69, no. 10, pp. 10506–10518, 2020.

[46] B. R. McAuliffe, M. Croken, M. Ahmadi-Baloutaki, and A. Raeesi, "Fuel-economy testing of a three-vehicle truck platooning system," tech. rep., National Research Council Canada. Aerodynamics Laboratory, Apr. 2017.

[47] C. Bonnet and H. Fritz, "Fuel Consumption Reduction in a Platoon: Experimental Results with two Electronically Coupled Trucks at Close Spacing," SAE Technical Paper 2000-01-3056, Aug. 2000.

[48] J. W. Ward, E. M. Stegner, M. A. Hoffman, and D. M. Bevly, "A Method of Optimal Control for Class 8 Vehicle Platoons Over Hilly Terrain," *Journal of Dynamic Systems, Measurement, and Control*, vol. 144, p. 011108, Jan. 2022.

[49] J. Ploeg, N. van de Wouw, and H. Nijmeijer, "Lp String Stability of Cascaded Systems: Application to Vehicle Platooning," *IEEE Transactions on Control Systems Technology*, vol. 22, no. 2, pp. 786–793, 2014.

[50] A. Fröberg, E. Hellström, and L. Nielsen, "Explicit Fuel Optimal Speed Profiles for Heavy Trucks on a Set of Topographic Road Profiles," SAE Technical Paper 2006-01-1071, 2006.

[51] S. Torabi, "Fuel-Efficient Truck Platooning using Speed Profile Optimization," Master's thesis, Chalmers University of Technology, Göteborg, Sweden, 2017.

[52] A. Alam, J. Mårtensson, and K. H. Johansson, "Experimental evaluation of decentralized cooperative cruise control for heavy-duty vehicle platooning," *Control Engineering Practice*, vol. 38, pp. 11–25, May 2015.

[53] A. A. Alam, A. Gattami, and K. H. Johansson, "An experimental study on the fuel reduction potential of heavy duty vehicle platooning," in *13th International IEEE Conference on Intelligent Transportation Systems*, (Funchal, Madeira Island, Portugal), pp. 306–311, IEEE, Sept. 2010.

[54] P. Smith, J. Ward, J. Pierce, D. Bevly, and R. Daily, "Experimental Results and Analysis of a Longitudinal Controlled Cooperative Adaptive Cruise Control (CACC) Truck Platoon," in *Proceedings of the ASME 2019 Dynamic Systems and Control Conference*, vol. 1, ASME, Oct. 2019.

[55] J. Siefert, E. Stegner, P. Snitzer, J. Ward, D. M. Bevly, M. Hoffman, and A. Kotz, "Using Demanded Power and RDE Aggressiveness Metrics to Analyze the Impact of CACC Aggressiveness on Heavy Duty Platooning Power Consumption," SAE Technical Paper 2021-01-0069, p. 10, Apr. 2021.

[56] H. Borhan, M. Lammert, K. Kelly, C. Zhang, N. Brady, C.-S. Yu, and J. Liu, "Advancing Platooning with ADAS Control Integration and Assessment Test Results," pp. 2021–01–0429, Apr. 2021.

[57] S. E. Li and H. Peng, "Strategies to minimize the fuel consumption of passenger cars during car-following scenarios," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, vol. 226, pp. 419–429, Mar. 2012.

[58] D. J. Chang and E. K. Morlok, "Vehicle Speed Profiles to Minimize Work and Fuel Consumption," *Journal of Transportation Engineering*, vol. 131, pp. 173–182, Mar. 2005.

[59] E. Wood, A. Duran, E. Burton, J. Gonder, and K. Kelly, "EPA GHG Certification of Medium- and Heavy-Duty Vehicles: Development of Road Grade Profiles Representative of US Controlled Access Highways," *SAE International Journal of Commercial Vehicles*, vol. 9, no. 2, p. 79, 2016.

[60] A. Alam, *Fuel-efficient heavy-duty vehicle platooning*. PhD Thesis, Electrical Engineering, KTH Royal Institute of Technology, Stockholm, Sweden, 2014.

[61] S. Xu, S. E. Li, B. Cheng, and K. Li, "Instantaneous Feedback Control for a Fuel-Prioritized Vehicle Cruising System on Highways With a Varying Slope," *IEEE Transactions on Intelligent Transportation Systems*, vol. 18, no. 5, pp. 1210–1220, 2017.

[62] J. Park, "Getting the Most Out of Cruise Control," May 2019. Accessed 2023-03-07 at https://www.truckinginfo.com/332181/getting-the-most-out-of-cruise-control.

[63] M. J. Droege, B. Black, S. Ashta, J. Foster, G. M. Shaver, N. Jain, and R. Thayer, "Heavy-duty truck platooning on hilly terrain highways: Methods for assessment and improvement," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, p. 095440702110676, Dec. 2021.

[64] I. Ibitayo, "Enhanced Class 8 Truck Platooning via Simultaneous Shifting and Model Predictive Control," Master's thesis, Purdue University, West Lafayette, Indiana, 2019.

[65] B. Black, "Optimization of Vehicle Dynamics for Enhanced Class 8 Truck Platooning," Master's thesis, Purdue University, West Lafayette, Indiana, 2020.

[66] L. Tsuei and Savaş, "Transient aerodynamics of vehicle platoons during in-line oscillations," *Journal of Wind Engineering and Industrial Aerodynamics*, vol. 89, pp. 1085–1111, Oct. 2001.

[67] T. Ard, B. Pattel, K. Fuhs, A. Vahidi, and H. Borhan, "Simulated and Experimental Verification of Fuel-Efficient Truck Platooning With Model Predictive Control Under Grade and Traffic Disturbances," *Journal of Autonomous Vehicles and Systems*, vol. 2, p. 031003, July 2022.

[68] A. Voronov, J. Andersson, and C. Englund, "Cut-ins in Truck Platoons: Modeling Loss of Fuel Savings," in *Towards Connected and Autonomous Vehicle Highways* (U. Z. A. Hamid and F. Al-Turjman, eds.), pp. 11–26, Cham: Springer International Publishing, 2021. Series Title: EAI/Springer Innovations in Communication and Computing.

[69] R. Sarkar, D. Verma, M. Hoffman, S. Lakshmanan, and B. Jakubowski, "Final Technical Report - Fuel-Efficient Platooning in Mixed Traffic Highway Environments," Tech. Rep. OSTI/1834543, Dec. 2021. Available at https://www.osti.gov/biblio/1834543.

[70] E. C. B. Olsen, S. E. Lee, W. W. Wierwille, and M. J. Goodman, "Analysis of Distribution, Frequency, and Duration of Naturalistic Lane Changes," *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, vol. 46, pp. 1789–1793, Sept. 2002.

[71] E. Nodine, A. Lam, M. Yanagisawa, and W. Najm, "Naturalistic Study of Truck Following Behavior," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2615, pp. 35–42, Jan. 2017.

[72] X. Wang, M. Yang, and D. Hurwitz, "Analysis of cut-in behavior based on naturalistic driving data," *Accident Analysis & Prevention*, vol. 124, pp. 127–137, Mar. 2019.

[73] P. Smith and D. Bevly, "Analysis of On-Road Highway Testing for a Two Truck Cooperative Adaptive Cruise Control (CACC) Platoon," SAE Technical Paper 2020-01-5009, Mar. 2020.

[74] A. Sarkar, J. Engstrom, and R. Hanowski, "Analysis of Car Cut-ins Between Trucks Based on Existing Naturalistic Driving Data," Tech. Rep. NSTSCE/22-UI-102, National Surface Transportation Safety Center for Excellence, Mar. 2022.

[75] Y. Dou, D. Ni, Z. Wang, J. Wang, and F. Yan, "Strategic car-following gap model considering the effect of cut-ins from adjacent lanes," *IET Intelligent Transport Systems*, vol. 10, pp. 658–665, Dec. 2016.

[76] F. Remmen, I. Cara, E. de Gelder, and D. Willemsen, "Cut-in Scenario Prediction for Automated Vehicles," in *2018 IEEE International Conference on Vehicular Electronics and Safety (ICVES)*, (Madrid), pp. 1–7, IEEE, Sept. 2018.

[77] S. P. Douglass, S. Martin, A. Jennings, H. Chen, and D. M. Bevly, "Deep Learned Multi-Modal Traffic Agent Predictions for Truck Platooning Cut-Ins," in *2020 IEEE/ION Position, Location and Navigation Symposium (PLANS)*, (Portland, OR, USA), pp. 688–697, IEEE, Apr. 2020.

[78] V. Milanés and S. E. Shladover, "Handling Cut-In Vehicles in Strings of Cooperative Adaptive Cruise Control Vehicles," *Journal of Intelligent Transportation Systems*, vol. 20, pp. 178–191, Mar. 2016.

[79] L. Johannesson, M. Nilsson, and N. Murgovski, "Look-ahead Vehicle Energy Management with Traffic Predictions," *IFAC-PapersOnLine*, vol. 48, no. 15, pp. 244–251, 2015.

[80] G. Guo and Q. Wang, "Fuel-Efficient En Route Speed Planning and Tracking Control of Truck Platoons," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, pp. 3091–3103, Aug. 2019.

[81] P. Bhagdikar, S. Gankov, S. Rengarajan, J. Sarlashkar, S. Hotz, and K. Bakshi, "Quantifying System Level Impact of Connected and Automated Vehicles in an Urban Corridor," SAE Technical Paper 2022-01-0153, Mar. 2022.

[82] W. Hucho and G. Sovran, "Aerodynamics of road vehicles," *Annual review of fluid mechanics*, vol. 25, no. 1, pp. 485–537, 1993.

[83] J. Howell, D. Forbes, and M. Passmore, "A drag coefficient for application to the WLTP driving cycle," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, vol. 231, pp. 1274–1286, Aug. 2017. Publisher: IMECHE.

[84] S. Windsor, "Real World Drag Coefficient – Is It Wind Averaged Drag?," in *The International Vehicle Aerodynamics Conference*, pp. 3–17, Woodhead Publishing, Nov. 2014.

[85] SAE J1252, "SAE Wind Tunnel Test Procedure for Trucks and Buses," Ground Vehicle Standard J1252_201207, SAE International, 2012.

[86] H. Chowdhury, H. Moria, A. Ali, I. Khan, F. Alam, and S. Watkins, "A Study on Aerodynamic Drag of a Semi-trailer Truck," *Procedia Engineering*, vol. 56, pp. 201–205, 2013.

[87] S. Wordley and J. W. Saunders, "On-road Turbulence," *SAE International Journal of Passenger Cars - Mechanical Systems*, vol. 1, pp. 341–360, Apr. 2008.

[88] S. Wordley and J. W. Saunders, "On-road Turbulence: Part 2," *SAE International Journal of Passenger Cars - Mechanical Systems*, vol. 2, pp. 111–137, Apr. 2009.

[89] B. R. McAuliffe, L. Belluz, and M. Belzile, "Measurement of the On-Road Turbulence Environment Experienced by Heavy Duty Vehicles," *SAE International Journal of Commercial Vehicles*, vol. 7, pp. 685–702, Sept. 2014.

[90] B. McAuliffe, H. Barber, and F. Ghorbanishohrat, "The Influence of Traffic Wakes on the Aerodynamic Performance of Heavy Duty Vehicles," SAE Technical Paper 2023-01-0919, Apr. 2023.

[91] B. McAuliffe and H. Barber, "Simulating Traffic-wake Effects in a Wind Tunnel," SAE Technical Paper 2023-01-0950, Apr. 2023.

[92] B. McAuliffe and M. Ahmadi-Baloutaki, "An Investigation of the Influence of Close-Proximity Traffic on the Aerodynamic Drag Experienced by Tractor-Trailer Combinations," SAE Technical Paper 2019-01-0648, Apr. 2019.

[93] A. A. Hussein and H. A. Rakha, "Vehicle Platooning Impact on Drag Coefficients and Energy/Fuel Saving Implications," *IEEE Transactions on Vehicular Technology*, vol. 71, pp. 1199–1208, Feb. 2022.

[94] H. A. Rakha, K. Ahn, K. Moran, B. Saerens, and E. V. d. Bulck, "Virginia Tech Comprehensive Power-Based Fuel Consumption Model: Model development and testing," *Transportation Research Part D: Transport and Environment*, vol. 16, pp. 492–503, Oct. 2011.

[95] B. McAuliffe, "A Drag-Reduction Prediction Model for Truck Platoons," SAE Technical Paper 2024-01-2548, Apr. 2024.

[96] E. Stegner, "Reverse Engineering the Correct Schmid Model Coefficients," Dec. 2023. Available at https://gitlab.com/ems0029/reverse-engineering-the-correct-schmid-drag-reduction-coefficients.

[97] F. Jaffar, T. Farid, M. Sajid, Y. Ayaz, and M. J. Khan, "Prediction of Drag Force on Vehicles in a Platoon Configuration Using Machine Learning," *IEEE Access*, vol. 8, pp. 201823–201834, 2020.

[98] Q. Luo, J. Li, and H. Zhang, "Drag coefficient modeling of heterogeneous connected platooning vehicles via BP neural network and PSO algorithm," *Neurocomputing*, vol. 484, pp. 117–127, May 2022.

[99] J. Dong, Q. Gao, J. Li, J. Li, Z. Hu, and Z. Liu, "Innovative modeling strategy of wind resistance for platoon vehicles based on real-time disturbance observation and parameter identification," *Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering*, p. 095440702311532, Feb. 2023.

[100] K. B. Devika, R. G., V. R. Shreya Yellapantula, and S. C. Subramanian, "A Dynamics-Based Adaptive String Stable Controller for Connected Heavy Road Vehicle Platoon Safety," *IEEE Access*, vol. 8, pp. 209886–209903, 2020.

[101] Q. Zhao, H. Zheng, C. Kaku, F. Cheng, and C. Zong, "Safety Spacing Control of Truck Platoon Based on Emergency Braking under Different Road Conditions," *SAE International Journal of Vehicle Dynamics, Stability, and NVH*, vol. 7, pp. 10–07–01–0005, Oct. 2022.

[102] A. T. Hamada and M. F. Orhan, "An overview of regenerative braking systems," *Journal of Energy Storage*, vol. 52, p. 105033, Aug. 2022.

[103] J. S. Chen, "Speed and Acceleration Filters/Estimators for Powertrain and Vehicle Controls," SAE Technical Paper 2007-01-1599, Apr. 2007.

[104] R. A. Rohrer, J. D. Luck, S. K. Pitla, and R. Hoy, "Evaluation of the Accuracy of Machine Reported CAN Data for Engine Torque and Speed," *Transactions of the ASABE*, vol. 61, no. 5, pp. 1547–1557, 2018.

[105] S. E. Marx, J. D. Luck, R. M. Hoy, S. K. Pitla, E. E. Blankenship, and M. J. Darr, "Validation of machine CAN bus J1939 fuel rate accuracy using Nebraska Tractor Test Laboratory fuel rate data," *Computers and Electronics in Agriculture*, vol. 118, pp. 179–185, Oct. 2015.

[106] J. L. Jimenez-Palacios, *Understanding and Quantifying Motor Vehicle Emissions with Vehicle Specific Power and TILDAS Remote Sensing*. Dissertation, Massachusetts Institute of Technology, Boston, MA, USA, Feb. 1999.

[107] H. Van Dop, "Terrain classification and derived meteorological parameters for inter-regional transport models," *Atmospheric Environment (1967)*, vol. 17, pp. 1099–1105, Jan. 1983.

[108] V. Girbés, D. Hernández, L. Armesto, J. Dols, and A. Sala, "Drive Force and Longitudinal Dynamics Estimation in Heavy-Duty Vehicles," *Sensors*, vol. 19, p. 3515, Aug. 2019.

[109] V. Girbes-Juan, L. Armesto, D. Hernandez-Ferrandiz, J. F. Dols, and A. Sala, "Asynchronous Sensor Fusion of GPS, IMU and CAN-Based Odometry for Heavy-Duty Vehicles," *IEEE Transactions on Vehicular Technology*, vol. 70, pp. 8617–8626, Sept. 2021.

[110] R. Stengel, *Optimal Control and Estimation*. Dover books on advanced mathematics, Dover Publications, 1994.

[111] T.A.S. Corporation, *Applied Optimal Estimation*. IT Pro, MIT Press, 1974.

[112] E. C. Ifeachor and B. W. Jervis, *Digital signal processing: a practical approach*. Pearson Education, 2002.

[113] North American Council for Freight Efficiency, "Trailer Aerodynamics," 2021. Available at https://nacfe.org/research/trailer-aerodynamics/. Accessed Feb. 2024.

[114] C. Walck, *Hand-book on statistical distributions for experimentalists*. Stockholm: Stockholms universitet, 1996.

[115] D. Cox and D. Oakes, *Analysis of Survival Data*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability, Taylor & Francis, 1984.

[116] F. S. Dias, M. Betancourt, P. M. Rodríguez-González, and L. Borda-de Água, "BetaBayes—A Bayesian Approach for Comparing Ecological Communities," *Diversity*, vol. 14, p. 858, Oct. 2022.

[117] G. James, D. Witten, T. Hastie, and R. Tibshirani, *An Introduction to Statistical Learning: with Applications in R*. Springer Texts in Statistics, New York, NY: Springer US, 2021.

[118] S. W. Menard, *Applied logistic regression analysis*. No. 106 in Quantitative applications in the social sciences, Thousand Oaks, Calif.: Sage Publ, 2. ed., [nachdr.] ed., 2008.

[119] Daimler Truck North America LLC, "eCascadia Specs | Freightliner Trucks." Available at https://www.freightliner.com/trucks/ecascadia/specifications/.

Appendices

Appendix A

Regression Analysis of the Available Methods

In this appendix, we reevaluate the suitability of the available methods to describe the energy consumption of a platoon. To do so, a regression analysis is performed on the entire Dataset 1, effectively quadrupling the sample size used in [10, 11]. First, Dataset 1 will be explored in terms of descriptive and marginal statistics. Then the model-building task begins, complete with variable selection and residual diagnostics. Once the best model candidate has been identified, the confidence intervals of the resulting model will be constructed. Finally, the model will be validated on unseen data for demonstration purposes.

A.1 Data Description

Dataset 1 is aggregated from 450 hour-long runs of test track platooning, conducted as part of the Department of Energy-funded contract DE-EE0008470. The total number of data points is 450, one for each aggregated test run. Each data point represents an hour of semi-truck operation, in a variety of configurations. For instance, there is data from:

- Four different semi trucks of varying masses and body styles
- Standalone driving, two-truck platooning, and four-truck platooning
- A variety of following distances between 35 and 100 feet
- A variety of following distance control strategies (fixed and flexible)
- Two different test tracks: the relatively flat NCAT test track in Opelika, AL, and the hilly ACM test track in Ypsilanti, MI.
- Four-truck tests where a passenger vehicle cuts in between the platoon multiple times per hour

The true goal of the analysis is to characterize the energy consumption of different platoons with a given set of regressors. Since the data lack a direct measurement of energy consumption, the reported fuel consumption of the vehicles is instead used as a proxy. Since

trucks have a vastly different fuel efficiency, fuel consumption is normalized using the standalone fuel consumption of each vehicle, already introduced as NFC. As in [10, 11], NFC is calculated by dividing the fuel consumed during platooning by the fuel consumed in isolation, as Figure A.1 shows. For example, NFC = 0.8 indicates that the truck consumed 80% as much fuel in the platoon as it did when driving alone, a 20% reduction in consumption. When there is zero drag reduction and braking, NFC is expected to be one.

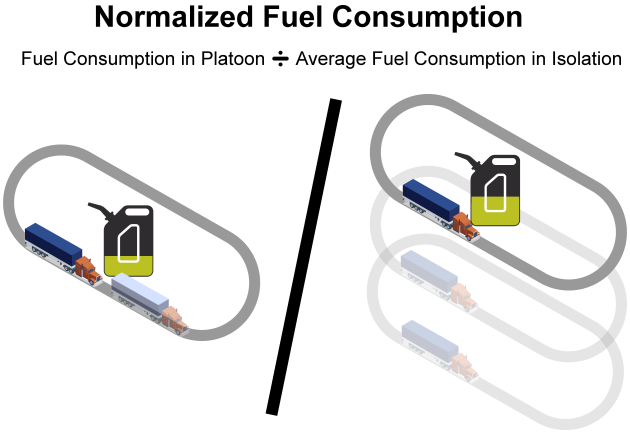

Figure A.1: Calculation process for Normalized Fuel Consumption (NFC), the basis for the target variable NFCr.

Some important insights are available by working with the NFC residual (NFCr), defined as

$$NFCr = NFC - \mathbb{E}\left[NFC \mid (E_{AD}, \text{Drag reduction} = 0)\right] = NFC - 1 \tag{A.1}$$

NFCr will be used as the target variable, since it facilitates a more natural test of the intercept hypothesis.

To describe NFCr, 14 regressors have been extracted from Dataset 1, 8 of which are numeric and 6 of which are indicators. Table A.1 presents the variables in the dataset.

Since the data consists of many different IVDs, it is expected that the aerodynamic drag reductions are vastly different across the data. To compensate for this, several candidate aerodynamic drag reduction models have been employed, including [93, 28][1]. The output of these models is a Drag Reduction Ratio (DRR), a unitless quantity that gives the reduction

[1]The model [28] has errors in the published values, and had to be reverse-engineered. See [96].

Table A.1: Dataset variables.

nfcr	:	Normalized Fuel Consumption residual (unitless)
ead	:	Total Active Deceleration Energy E_{AD} (kJ/kg·hr)
drf$_1$	:	Drag Reduction Fraction #1 by Hussein, power law [93] (unitless)
drf$_2$	:	Drag Reduction Fraction #2 by Hussein, rational polynomial [93](unitless)
drf$_3$	:	Drag Reduction Fraction #3 by Schmid, reduced [28](unitless)
drf$_4$	:	Drag Reduction Fraction #4 by Schmid, full model [28] (unitless)
rho	:	Air density from OpenWeatherMap Model (kg/m^3)
fan_pwr	:	Semi-truck estimated fan power (kW)
fan_uptime	:	Duty cycle for fan on $\in [0, 1]$
two_truck	:	Indicator variable for two-truck platooning (true/false)
four_truck	:	Indicator variable for four-truck platooning (true/false)
A2	:	Indicator variable for truck model A2 (true/false)
T13	:	Indicator variable for truck model T13 (true/false)
T14	:	Indicator variable for truck model T14 (true/false)
leading	:	Indicator variable for whether the truck led the platoon (true/false)

in drag as a ratio of platoon drag to baseline drag. A value of 0.8 here would similarly represent a 20% drag reduction versus standalone operation.

It is more interpretable to represent the effect of drag reduction as a percentage reduction (drag reduction fraction: $DRF = 1 - DRR$). For example, a DRF of 0.2 represents a 20% reduction in drag, and DRF = 0 means no change to aerodynamic drag. Without braking, a good drag reduction model should be able to describe the energy savings from platooning well.

Another regressor is the active deceleration energy E_{AD}, which is calculated according to the method in [11]. It is a calculated value that uses a predictive vehicle model and therefore depends on the quality of the model. Figure A.2 shows the mechanism used to calculate E_{AD} every time active deceleration is detected. The slight downward trend in the difference at the 100 second mark of Figure A.2 is probably due to modeling error. Units of E_{AD} are kilojoules per kilogram per hour, representing the amount of energy lost per hour for each kilogram of the vehicle. This enables semi-trucks of different masses to be analyzed together, which is the case for this dataset.

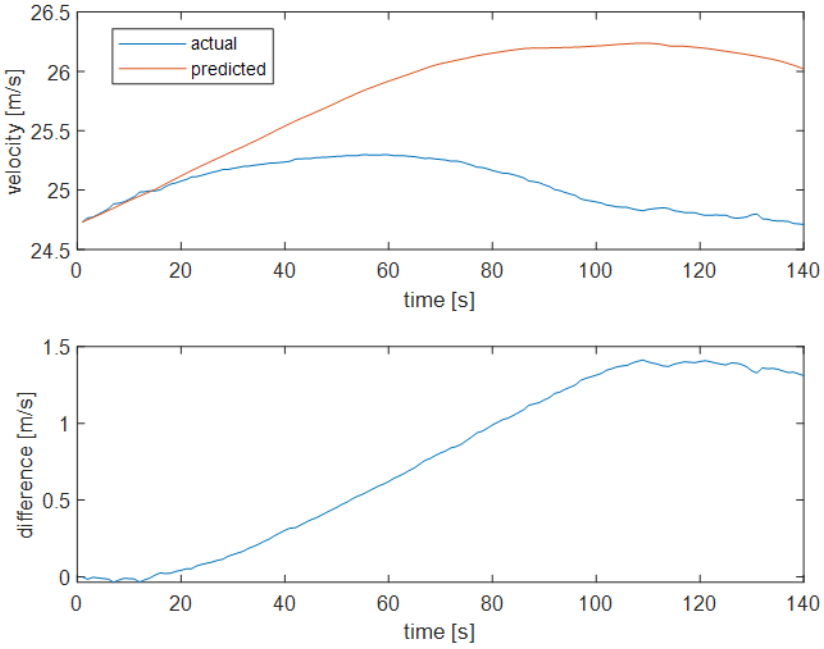

Figure A.2: A demonstration of the principle behind E_{AD} calculation, showing an instance of braking whereby a truck lost energy.

A semi-truck engine fan can consume up to 30 kilowatts (40 hp) of power [22], leading to a measurable increase in fuel consumption. As such, an engine fan power estimate has been included in the model, in units of kilowatts. Additionally, the engine fan duty cycle (on/off) has been included.

Many indicator variables have also been included. The platoon type has been one-hot encoded, which yields a dummy variable for two- and four-truck platoon operation. When neither the two- or four-truck indicator variables are active, it implies standalone operation. There are also dummy variables for each of the trucks, identified as "A2", "T13", and "T14". The indicator variables truly represent two categorical variables: the number of trucks in the platoon, at 3 levels, and the type of truck at four levels. This combines a total of 12 combinations of trucks and the number of trucks in the platoon. Lastly, an indicator of whether the truck led the platoon is included.

A.2 Exploratory Data Analysis of Dataset 1

The first task of data analysis is to examine the composition of the variables. As shown in Figure A.3, the dataset is relatively balanced by platoon and truck type, although there are fewer observations for the truck "T14" due to a transmission failure in the middle of the testing campaign.

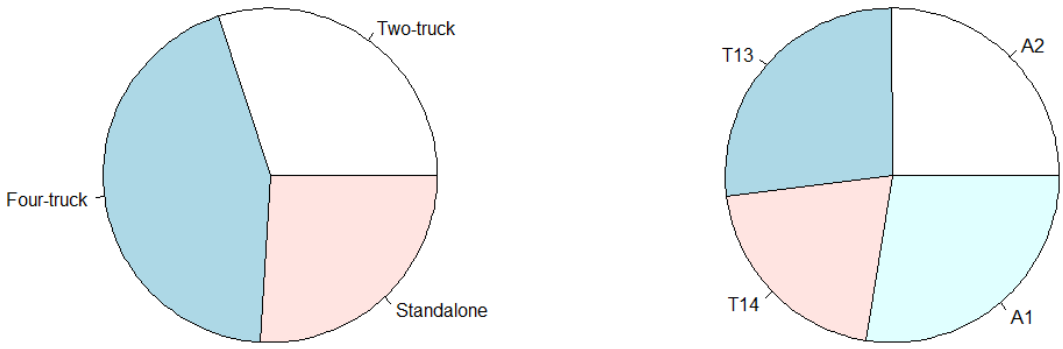

(a) Number of observations by platoon type. (b) Number of observations by truck type.

Figure A.3: Pie charts of the data composition.

A table of summary statistics helps to understand the characteristics of the dataset. Based on the information in Table A.2, we see that:

- The average NFCr is nearly 0 for the whole dataset, but goes up to 50% higher, and 19% lower than the standalone references.
- E_{AD} is effectively zero for at 50% of the dataset, but rises sharply between the 75th and 100th percentile
- The four distinct drag reduction ratios provide similar values
- There is little variation in air density

- The engine fan was on 35% of the time, with an estimated 4 kilowatt draw on average.

Table A.2: Summary statistics of the dataset.

Variable	Type	Counts	Mean	Std.	0th %	25th %	50th %	75th %	100th %
nfcr	numeric	-	-0.012	0.088	-0.191	-0.063	-0.012	0.021	0.496
ead	numeric	-	0.496	1.195	-0.001	0.000	0.008	0.266	8.574
drf1	numeric	-	0.097	0.092	0.000	0.000	0.080	0.172	0.280
drf2	numeric	-	0.105	0.100	-0.004	0.000	0.079	0.201	0.271
drf3	numeric	-	0.086	0.087	0.000	0.000	0.033	0.167	0.241
drf4	numeric	-	0.087	0.087	0.000	0.000	0.033	0.169	0.241
rho	numeric	-	1.207	0.036	1.141	1.180	1.208	1.239	1.280
fan_pwr	numeric	-	3.999	5.175	0.000	0.000	0.001	9.372	18.033
fan_uptime	numeric	-	0.350	0.457	0.000	0.000	0.000	1.000	1.000
two_truck	factor	135	-	-	-	-	-	-	-
four_truck	factor	198	-	-	-	-	-	-	-
A2	factor	113	-	-	-	-	-	-	-
T13	factor	121	-	-	-	-	-	-	-
T14	factor	92	-	-	-	-	-	-	-
leading	factor	146	-	-	-	-	-	-	-

Figure A.4 provides an in-depth look at how NFCr varies with E_{AD}. Per Figure A.4a, the standalone runs are centered about NFCr= 0, as they should be. The second notable behavior is that NFCr has a much lower value at $E_{AD} = 0$ if the truck was platooning. This is due to drag reduction while platooning, which can decrease fuel consumption considerably. The scatter in the standalone run of NFCr is normally distributed, and has a standard deviation of about 3%; see the histogram in Figure A.5.

The two-truck runs cover a smaller range of NFCr and E_{AD} than the four-truck runs, but there is no indication that the trend is different.

Figure A.4b shows that the truck A2 produced the points with the highest E_{AD}. This is because it occupied the last position in the four-truck platoons, thereby dealing with the worst braking requirements. Trucks T13 and T14 exhibited a measurable level of E_{AD}, where truck A1 has just one example of high E_{AD}[2].

[2]Truck A1 could not serve as a platoon follower, so it always led the platoon. This explains why it did not need to brake very often.

As for Figures A.4c and A.4d, the drag reduction fractions appear to correlate well, given $E_{AD} = 0$. Comparison of figure A.4c with A.4d shows that DRF_2 is more bimodal than DRF_4 and predicts a higher drag reduction in general. Trucks that led the platoon have more modest DRF's as shown by the colors in Figure A.4d. Because the range of DRF is small for leaders, it may be spurious to fit a separate DRF coefficient for leading trucks.

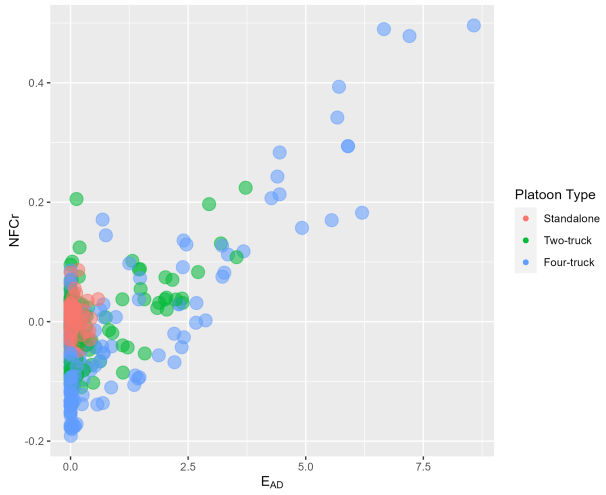

(a) NFCr versus E_{AD}, colored by platoon type.

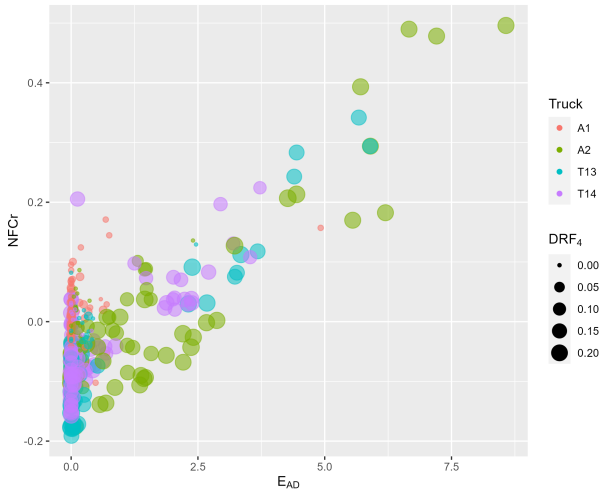

(b) NFCr versus E_{AD}, colored by truck.

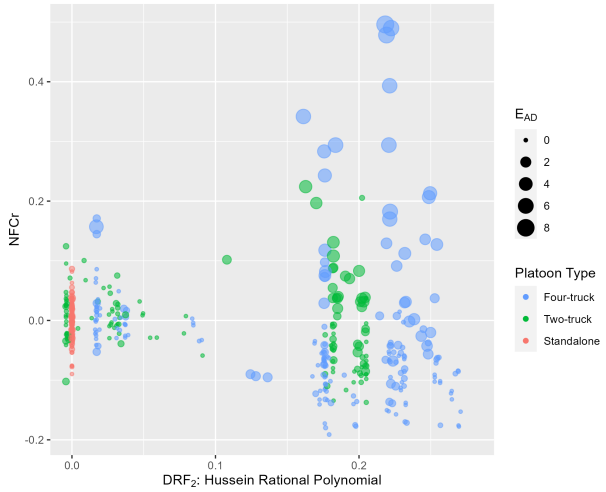

(c) NFCr versus DRF_2, colored by platoon type.

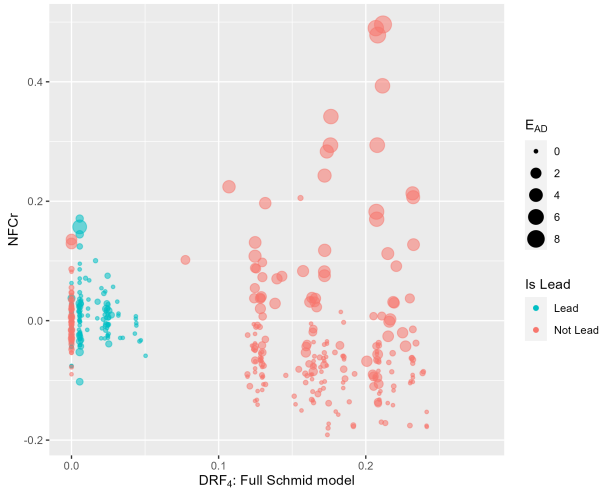

(d) NFCr versus DRF_4, colored by whether the truck led a platoon or not.

Figure A.4: NFC residual versus various regressors.

Based on this exploratory analysis, E_{AD} and some DRF's show good potential as predictors of NFCr. Next, the model building will be covered.

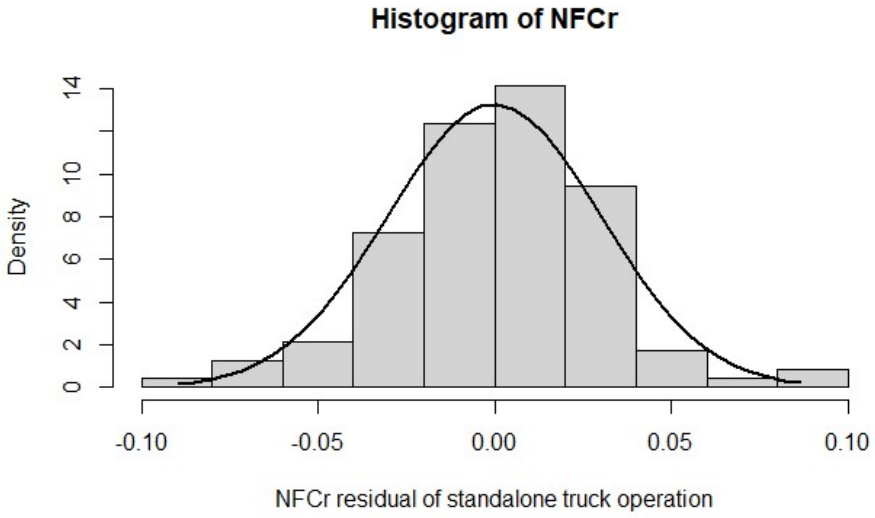

Figure A.5: Histogram of the NFCr for standalone trucks only.

A.3 Numeric Variable Selection by LASSO

There are many ways to build a linear regression model. In this case, multiple variables fill the same role: four drag-reduction submodels and two fan submodels. Selection of the best submodels should be carried out.

One increasingly popular method to select variables is regularized LASSO regression combined with cross-validation, which splits a dataset into folds, holding one fold out at a time, while the rest are used to fit the model. LASSO regularization is excellent at cutting variables from models as the penalty on coefficient size increases. Here, LASSO regression was used to downselect the DRF and fan models. The process used is as follows:

1. First, the dataset was split into 10 randomized folds.
2. Next, each regressor was normalized by centering and scaling to avoid discriminant penalties
3. LASSO regressions with shrinkage factors from 10^{-10} to 1 were fit, with the value of the coefficients plotted versus shrinkage in Figure A.6.

4. Variables whose coefficients decreased to zero were removed from the running, even if the size of the coefficient was not monotonic with penalty λ.

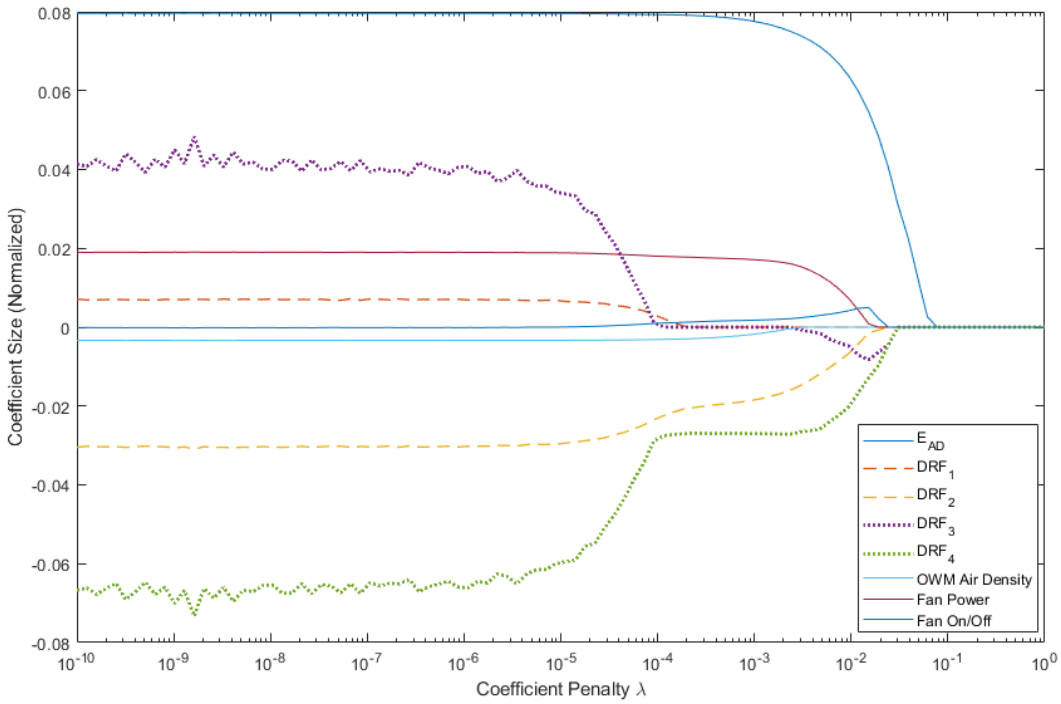

Figure A.6: LASSO regularization to select variables.

E_{AD} survives the longest of all variables, confirming its value to the model. It was determined that DRF_4 (the Schmid model was the best of the drag reduction submodels based on its longevity as well. The fan power estimate was deemed better than the fan on/off model. Air density does not survive very long and does not merit further mention.

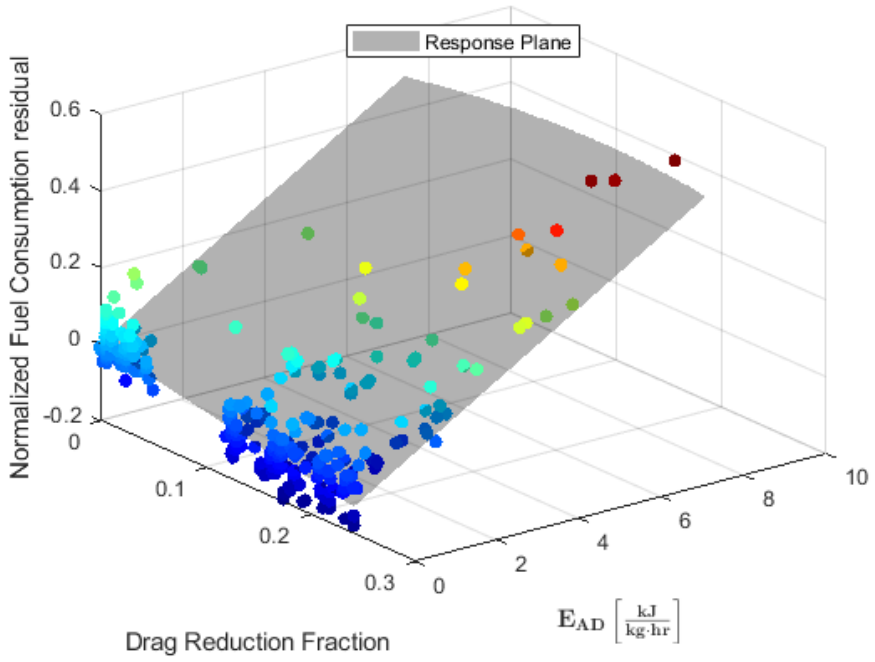

Figure A.7: Two-term model response plane.

A.4 The Two-Term Model

At a minimum, the final model will need to use E_{AD} and drag reduction term, since these are dominant mechanisms in platooning energy use. Based on this fact, a two-term linear model of the following form was constructed:

$$NFCr = \beta_0 + \beta_1 E_{AD} + \beta_2 DRF_4 + \varepsilon \tag{A.2}$$

$$\hat{NFCr} = 0.0018 + 0.0659 E_{AD} - 0.541 DRF_4 \tag{A.3}$$

Figure A.7 shows the response plane for the two-term model. For the model, the R^2 value was 0.75[3], indicating that 75% of the variability in NFCr is explained by regression on E_{AD} and DRF_4.

[3]This R^2 value and subsequent ones were calculated after removing observation 161, which was deemed invalid for reasons to be discussed later.

A.5 Intercept Hypothesis Testing

It is reasonable to believe that the two-term model should have no intercept, since NFCr should be zero in the absence of braking/drag reduction. If it does not equal zero, then that may indicate that NFCr is biased. What statistical evidence is there for a no-intercept model?

The hypothesis to be tested is:

$$H_o : \beta_0 = 0$$

$$H_a : \beta_0 \neq 0$$

The p-value for β_0 with the two-term model in Equation A.2 is 0.531, which is insignificant at any reasonable confidence interval. However, removing the intercepts from linear models is rarely advisable, and leaving them in does no harm. If the intercept is not zero, then an adjustment to NFCr will have to be made to continue interpreting it as a percent change in fuel consumed. This adjustment is made because NFCr should be centered at zero given that the drag reduction and E_{AD} are zero, and it should scale with NFC = 1. Equation A.4 shows the adjustment to NFCr to maintain the interpretation that NFCr represents *the change in fuel consumption as ratio.*

$$N\hat{F}Cr_{adj} = \frac{N\hat{F}Cr - \hat{\beta}_0}{1 + \hat{\beta}_0} \tag{A.4}$$

A.6 Adding More Terms to the Model

For a three-term model, the results of LASSO pointed to the addition of the variable *fan power*. Adding the fan power estimate increases the variability explained from 75% to 80%. This increase in R^2 is understandable considering Figure A.8. Many of the points that previously sat above the plane have been pushed downward towards it, effectively pinching it

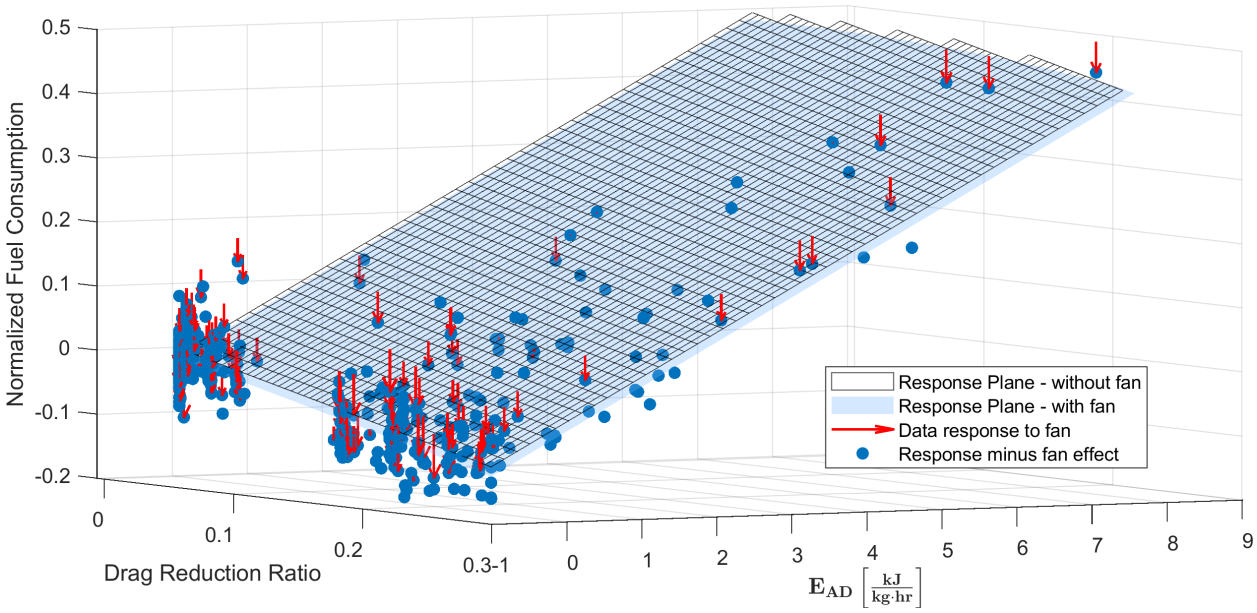

Figure A.8: Change in the response due to the addition of fan power.

in place from the top. As a result, the response plane is translated downward, which makes the intercept term significant.

Beyond the three-term model, the six indicator variables in the dataset were investigated. In particular, the interactions between drag reduction and different trucks, platoon types, and positioning were checked. A complete model was built and the R `step` function using BIC was used to select a model. The suggested model included *leading* and an interaction between *leading* and E_{AD}:

$$NFCr = \beta_0 + \beta_1 E_{AD} + \beta_2 DRF_4 + \beta_3(\text{fan power}) + \beta_4(\text{leading}) + \beta_5 E_{AD} \times (\text{leading}) + \varepsilon \quad \text{(A.5)}$$

The additions to the five-term model over the three-term model are probably due to the small range of *leading* data, as shown in Figure A.4d. Although R^2_{adj} did increase to 0.814, the small range of E_{AD} in the *leading* data led to choosing the three-term model over the five-term model.

A.7 Final Model Selection

Table A.3: Statistics for various candidate models with and without fan power.

Model	R^2, adj	Intercept β_0	p-value for β_0
Two-term, Eq. A.2	0.7486	0.0018	0.53
Three-term, Eq. A.6	0.8005	-0.0142	$3E-06$
Five-term, Eq. A.5	0.8139	-0.0197	$8E-07$

Table A.3 shows the two-, three- and five-term model summary statistics. Ultimately, the three-term model was selected as the best candidate. Equation A.6 shows the form of the model and its coefficients:

$$NFCr = \beta_0 + \beta_1 E_{AD} + \beta_2 DRF_4 + \beta_3(\text{fan power}) + \varepsilon \tag{A.6}$$

$$\hat{NFCr} = -0.014 + 0.066 E_{AD} + -0.53 DRF_4 + 0.0039(\text{fan power}) \tag{A.7}$$

A quick interpretation of the coefficients in Equation A.7 is provided in Table A.4.

Table A.4: Interpretation of model coefficients and their 95% confidence intervals, re-normalized per Equation A.4.

Coefficient	Units	Value	Interpretation: Effect on Fuel Consumption
$\hat{\beta}_1$	$^1/_{\frac{\mathbf{kJ}}{\mathbf{kg \cdot hr}}}$	0.065956	1 $\frac{\mathbf{kJ}}{\mathbf{kg \cdot hr}}$ of E_{AD} yields a $6.7 \pm 0.6\%$ increase
$\hat{\beta}_2$	$^{100}/_{\text{drag reduction \%}}$	-0.533625	A 10% drag reduction yields a $5.4 \pm 0.5\%$ decrease
$\hat{\beta}_3$	$^1/_{\mathbf{kW}}$	0.003865	10 **kW** of fan power yields a $3.9 \pm 0.7\%$ increase

To summarize, platooning is meant to save fuel, but in real-world environments, the savings are often compromised. A model has been proposed to estimate the change in energy consumption due to platooning. The selected model is a function of E_{AD}, the drag reduction model in [28], and an estimate of the engine fan power. Now that the rationale for the final model has been provided, the attention will turn to a model adequacy check.

A.8 Normality of Residuals and Influential Points

Ordinary least squares regression makes several assumptions. One of the key assumptions is that the residuals are normally distributed. Often, this assumption of normality is checked by using a Q-Q plot, which places the sample quantiles versus the externally studentized residual quantiles. Figure A.9a shows the Q-Q plot for the selected model, and Figure A.9b shows the residual histogram. It appears that the residuals are reasonably normal. The departure from normality is not excessive. Observations 161 and 15 were flagged for further investigation due to their excessive residuals.

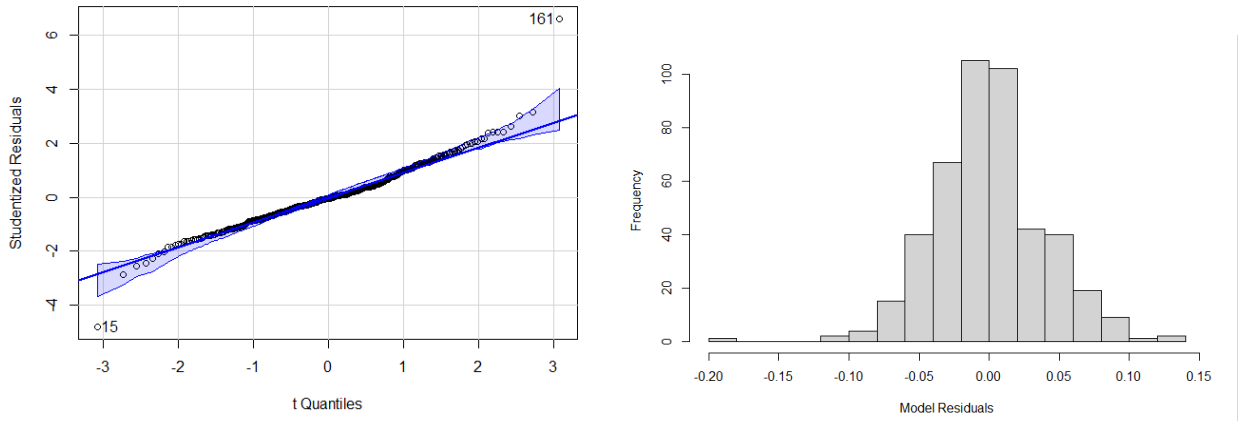

(a) Q-Q-plot of the model residuals, highlighting observations 15 and 161 as outliers.

(b) Histogram of the model residuals.

Figure A.9: Plots to check the normality of the least-squares model assumptions.

Observation 15 is the only observation with low drag reduction but high E_{AD} (Figure A.10a). It was kept, but in such cases, the model may be inadequate due to the sparsity of points in the region. On the other hand, Observation 161 was discarded, as it showed unrealistic velocity behavior (Figure A.10b), and the test notes showed that the truck was operating poorly.

It appears that the assumptions of the model are reasonably satisfied.

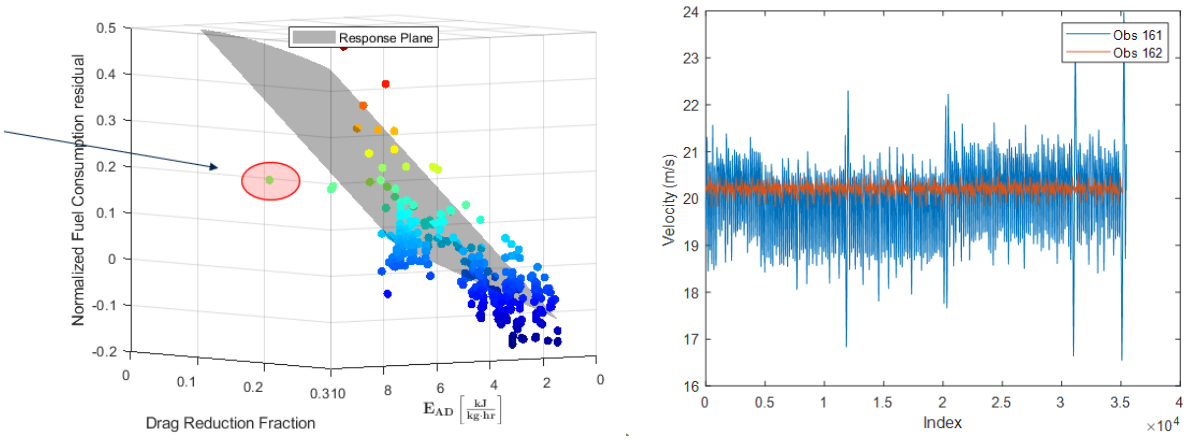

(a) The location of observation 15, an outlier.

(b) The velocity trace for observation 161 plotted along with a velocity trace for a successful trial.

Figure A.10: Supporting figures for the outliers.

Appendix B

The Effects of Net Elevation Change on Platoon Energy Consumption

It is unrealistic to assume that platooning microtrips will have no net elevation change. To illustrate the effect of including a net elevation change in platoon energy calculations, consider the following scenario in which a platoon experiences a reduction in drag, braking losses, and a net elevation change.

First, imagine that a truck is operating on level ground at a constant speed. This truck's mean reference power consumption is:

$$\bar{P}_{ref} = \bar{P}_{aero,ref} + \bar{P}_{etc,ref} \tag{B.1}$$

where $\bar{P}_{aero,ref}$ is mean aerodynamic drag power and $\bar{P}_{etc,ref}$ is encompasses everything else, namely rolling resistance and accessory loads.

Now imagine that this truck becomes part of a platoon and achieves a 20% reduction in drag by platooning, that is, $DRR = 0.8$. For simplicity, we will declare that the power generated by the truck is divided equally between the aerodynamic drag and other forces experienced on the road. Thus, 50% of its power goes to aerodynamic drag and the other 50% to rolling resistance and accessory loads. According to these parameters, when the truck is on level ground and the brakes are not applied, there would be a 10% decrease in energy consumption for this particular truck.

Now consider the case where the act of platooning has caused the truck to brake, leading to an increase in energy usage. We will denote this braking load as a percentage of the combined aerodynamic drag and other road loads using the symbol β:

$$\begin{aligned} \bar{P}_{AD} &= \beta(\bar{P}_{ref}) \\ \beta &\geq 0 \end{aligned} \tag{B.2}$$

This representation is equivalent to the percentage by which power consumption increases due to braking. If $\beta = 10\%$, i.e. the braking load is 10% of the reference drag and other road loads, it would effectively counterbalance the 10% decrease in energy consumption resulting from the reduction of aerodynamic drag.

Now finally, imagine that the platooning truck also encounters a net elevation change. Unlike the braking load, the gravitational load can be negative or positive. We choose to express the gravitational load as a fraction (γ) of the aerodynamic drag and other road loads.

$$\bar{P}_{grvt} = \gamma(\bar{P}_{ref}) \tag{B.3}$$

In mathematical terms, the NPC of the truck in the scenario we have described is:

$$\begin{aligned} NPC &= \frac{DRR \times \bar{P}_{aero,ref} + \bar{P}_{etc,ref} + \bar{P}_{AD} + \bar{P}_{grvt}}{\bar{P}_{aero,ref} + \bar{P}_{etc,ref} + \bar{P}_{grvt}} \\ &= \frac{\bar{P}_{aero,ref}(DRR + \gamma + \beta) + \bar{P}_{etc,ref}(1 + \gamma + \beta)}{\bar{P}_{ref}(1 + \gamma)} \end{aligned} \tag{B.4}$$

Note that when the gravitational load is equal and opposite to the reduction in drag and the braking losses, Equation B.4 becomes undefined (γ equals -100%). Furthermore, the limit as γ approaches -100% goes to infinity:

$$\lim_{\gamma \to -100\%^+} NPC = -\infty$$

Ultimately, the form of Equation B.4 allows us to visualize how braking losses and elevation will interact with NPC. Figure B.1 shows the NPC as a function of β and γ. As expected, there is a 10% savings when $\gamma = \beta = 0$. When $\beta = 10\%$, the benefit of drag reduction is canceled by braking losses, making NPC exactly one.

Positive elevation changes serve to dilute the other terms, as the truck spends an increasing proportion of its power on drag reduction. For example, at $\gamma = 100\%$, the energy savings of the platoon are half what they were without elevation change, 5% instead of 10%.

For negative elevation changes, NPC decreases exponentially as the negative change grows larger. This is because the drag savings remain constant, but the overall load decreases. As the net elevation change becomes increasingly negative, NPC approaches negative infinity. For a large enough elevation decrease, the platooning truck would have no load at all and can coast, while a non-platooning truck would still be using some power, which explains why the limit of NPC goes to negative infinity.

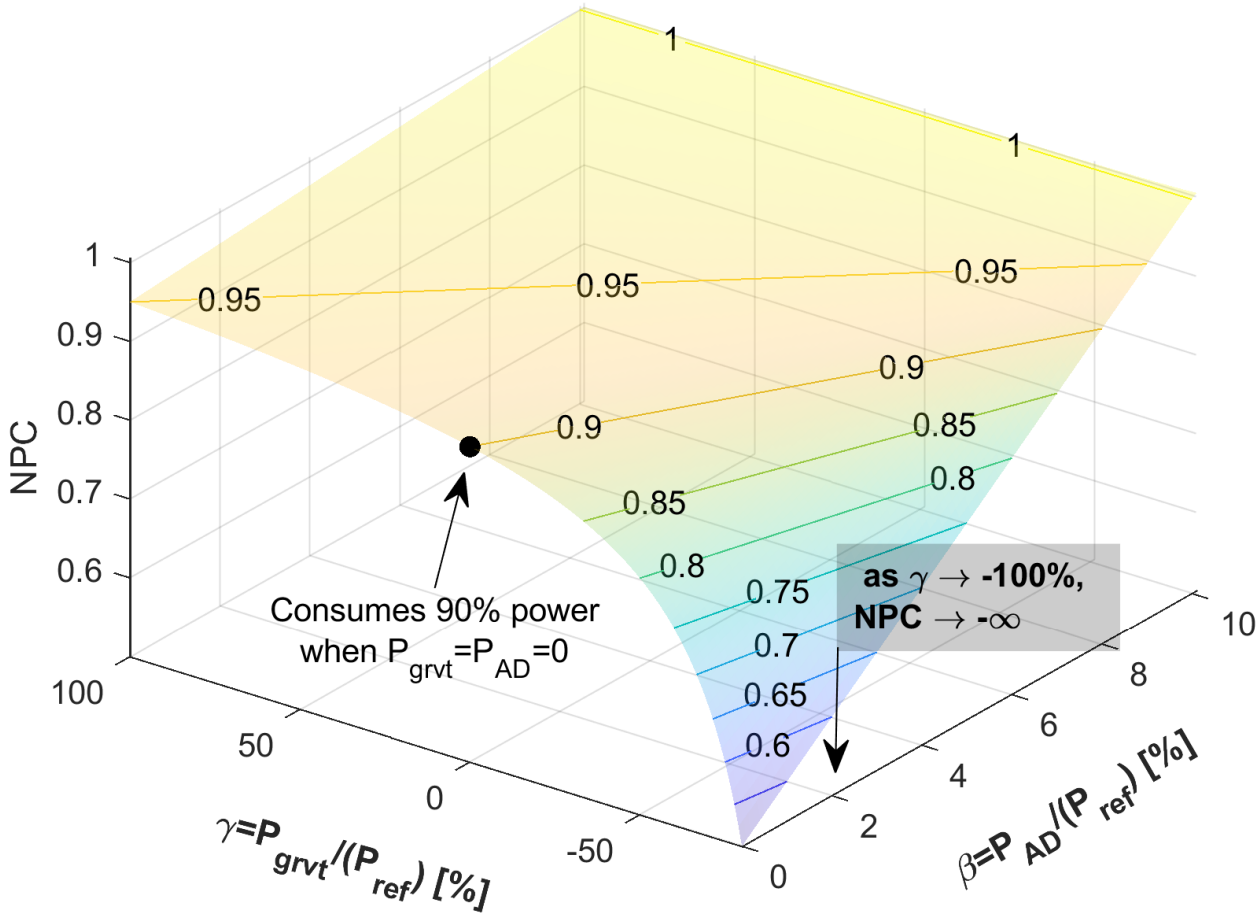

Figure B.1: Gravity and braking effects on NPC for a fixed drag reduction of 20% and an even power split between drag and other loads.

In summary, platooning on a net positive elevation will diminish the impact of platooning. It is crucial to note that both the reduction in drag and the losses in braking during platooning are weakened when driving uphill. This is because the additional load to overcome gravity will dilute the drag reduction savings as well as the braking losses.

As for platooning on a net negative elevation, the percent benefits of platooning actually increase exponentially. However, this can be misleading for two reasons. First, savings increase exponentially on a relative basis (e.g. %), not absolute (e.g. $, gallons, kWhr). Less total energy is consumed on downhill segments, so while a 40% reduction in fuel consumed appears to be a very large reduction, if the non-platooning energy use was already modest, the absolute energy savings are also modest.

Furthermore, it is probable that downhill sections require more braking, which will diminish the benefits of reduced drag [60]. To fully exploit the opportunity to save energy by platooning downhill, a lookahead controller is necessary, and the hills must have limited length and gradient to avoid speed limit-enforced braking.

Appendix C

Simulation Details

A heavy-duty vehicle simulation was developed in TruckMaker 11.1.1, a high-fidelity virtual test track software developed by IPG Automotive GmbH. The purpose of this appendix is to document the creation of this simulation environment, which was used throughout the dissertation to verify the construction of the framework, most notably in the model formulation and practical calculation sections of Chapter 3 (3.3 and 3.4), and later in the sensitivity analysis of Chapter 4 (4.1).

C.1 Environment

The simulation environment used in this dissertation was previously documented in [11], and some portions of this appendix are direct from that work, to which the author has copyright. The simulation was designed to replicate a realistic convoy of trucks navigating hilly terrain. The simulation parameters were adjusted to align with the experimental results of Dataset 1.

To create the virtual test track, the measured road grade profile of the 3702 m ACM highway loop was imported into TruckMaker. Unlike the ACM track, which is a loop, the virtual track was constructed without turns and included 1 km runways at the beginning and end of the elevation profile. Figure C.1 shows the elevation and grade profile of both the simulation environment and the measured test track data. It can be seen that the elevation and grade profiles are reasonably well matched to the measured elevation and grade. The purpose of the 1 km flat section is to allow the initial vehicle dynamics to stabilize prior to operation on the elevation profile.

Although there is a capability in TruckMaker to model additional traffic, wind, and other factors, in this simulation the emphasis was on creating realistic controller-induced braking losses, which was achieved merely by platooning over the hills.

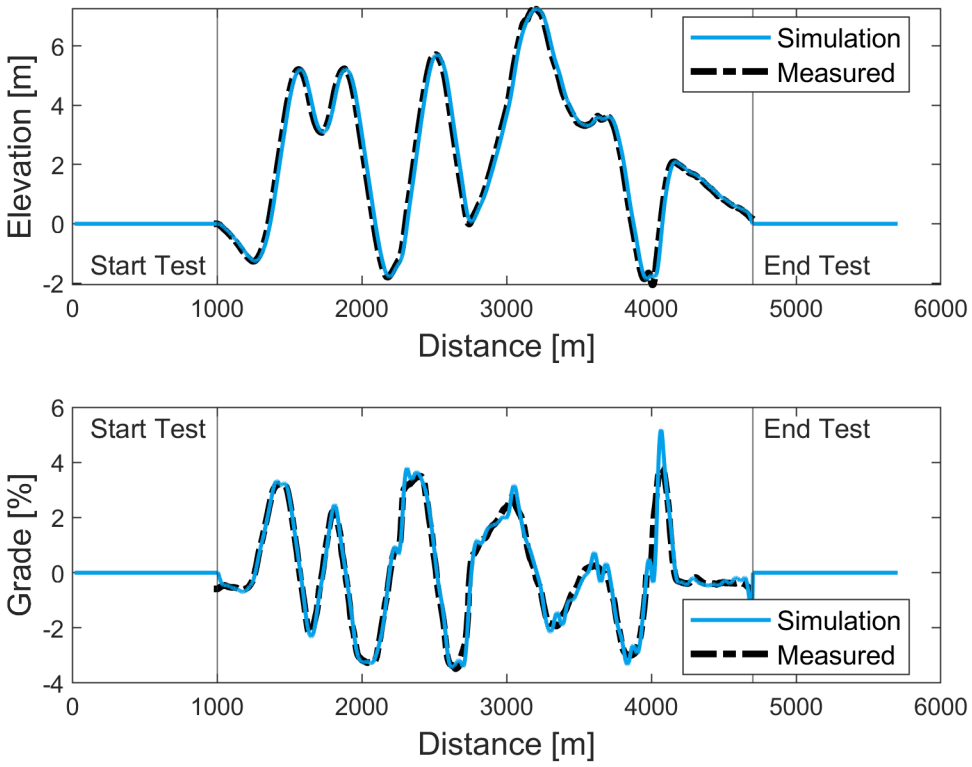

Figure C.1: The elevation and grade profiles of the simulation environment. The elevation and grade profiles from the experimental test track are overlaid.

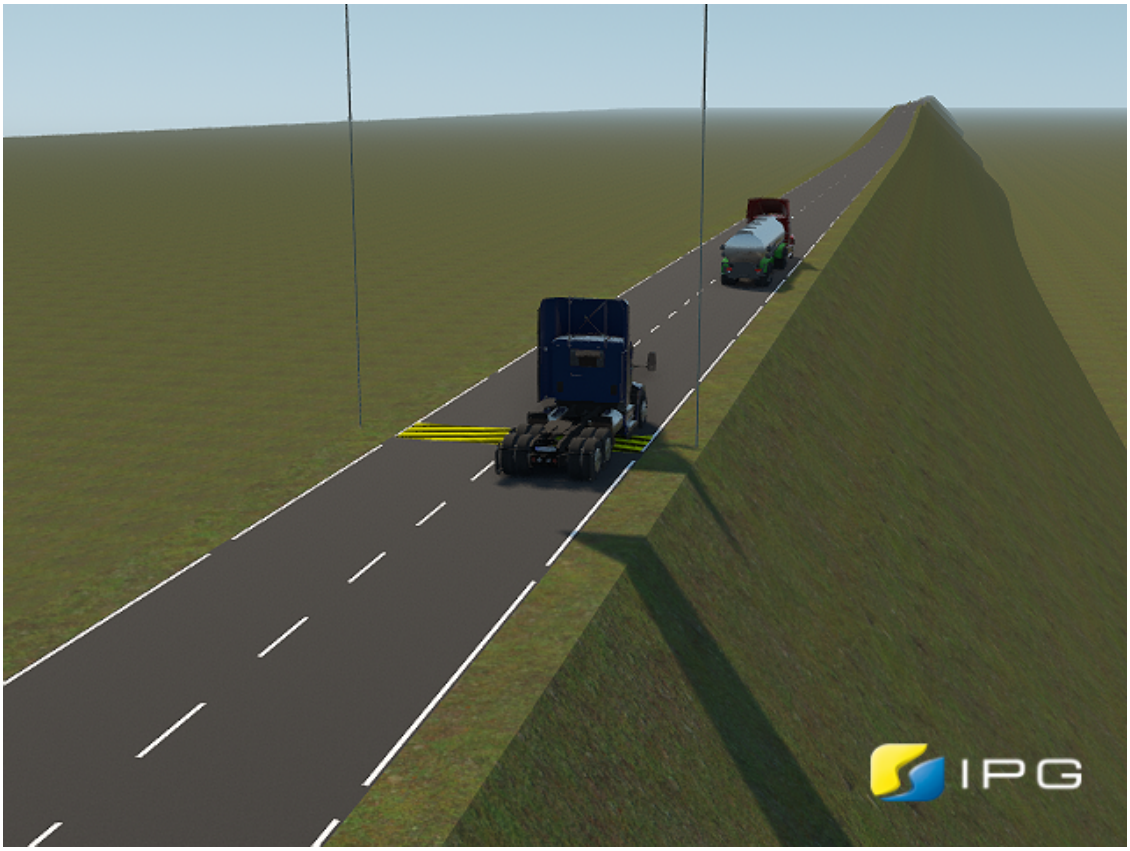

Figure C.2: Screenshot of the vehicle simulation as it passes the 1 km mark and begins the simulated grade.

C.2 Vehicle Model

Within TruckMaker, only one vehicle is fully simulated at a time (the ego vehicle). All other simulated vehicles are treated as traffic objects and follow kinematic commands only. Traffic objects act as points of reference for the subject vehicle and do not impact the aerodynamics of the subject vehicle. Figure C.2 shows a visualization of the simulation as the ego vehicle passes the 1 km mark where the ACM grade begins.

The simulated vehicle itself is based on an included model of a MY2021 Kenworth T680 daycab. The weight of the truck is set to 8000 kg, and a load of 7000 kg is placed at the approximate center of the vehicle to represent an empty trailer's mass and bring the base vehicle weight to 15000 kg. Figure C.3 shows a screenshot of the additional mass on the truck, with the empty 7000 kg trailer mass shown as a red dot. The simulated trailer load

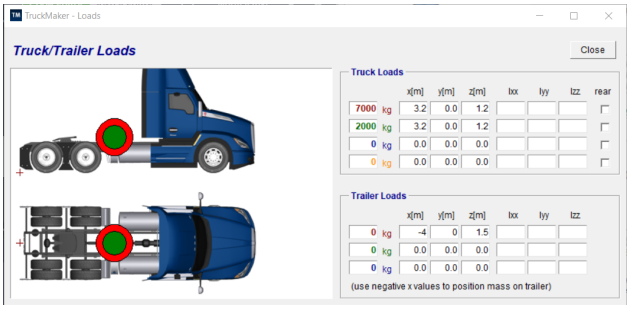

Figure C.3: Screenshot of the simulated trailer loads on the vehicle.

(which varied between simulations) was placed at the same location, as shown by the green 2000 kg dot in Figure C.3. The capability to model actual articulated trailers does exist in TruckMaker, this would unnecessarily complicate the post-calculation of braking losses with little gain of fidelity.

The simulated suspension and steering of the truck were left unmodified, with no modeled wheel bearing or suspension deflection losses. The tires of the truck were generated using the IPG Tire Data Set Generator. Their rolling resistance was set to 0.0098, without a velocity-dependent component. The brake system of the simulated truck was set as a hydraulic system to speed up simulations, rather than a more realistic air brake system. To approximate the longer delay in air brake actuation, the response time and the build-up time of 0- 75% brake torque were set to 0.2 seconds. Both of these parameters are documented in the TruckMaker documentation.

The powertrain model was given a higher level of customization than some of the other subsystems, due to the emphasis on braking power. The torque curve of the truck A1 was used. The turbocharger dynamics were neglected. The negative load at zero torque was set to essentially zero, to avoid confusing the error in the estimated braking power with the parasitic losses of the engine. A load- and speed-dependent specific fuel consumption map was implemented, which includes the realistic high-load "sweet spot" of modern 12-15 L turbodiesel engines. A screenshot of this fuel consumption map is shown in Figure C.4, which also shows the full-load torque curve of the engine.

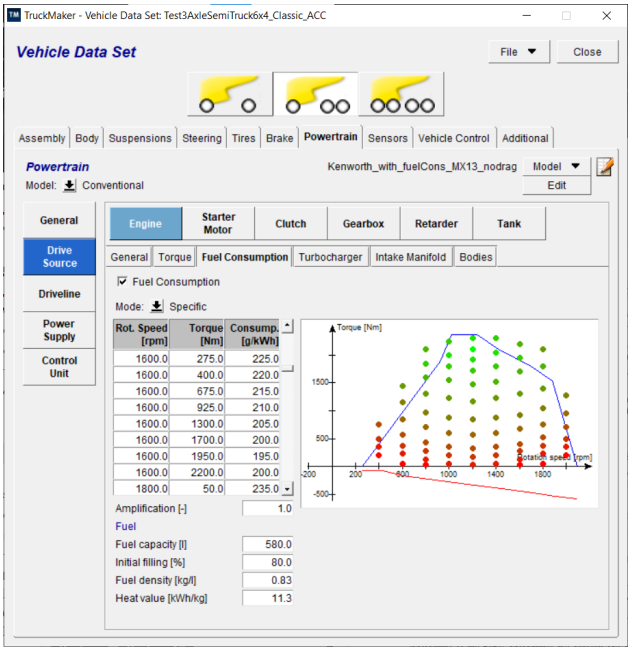

Figure C.4: Screenshot of the fuel consumption map settings for the powertrain of the simulated truck.

The vehicle was configured with a manual transmission, with the shift point set to 1800 RPM. The gear ratios and the final drive were set to match the truck A1, which has an Eaton Fuller automated manual transmission. Table C.1 shows the gear ratios of the simulated vehicle.

Table C.1: Gear ratios of the simulated truck, to mimic truck A1.

Final Drive	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th
2.579	11.06	10.20	7.06	4.98	3.96	2.83	2.03	1.47	1.00	0.79

The simulated IVD control strategy is described in the main body of the dissertation, as it was deemed relevant to the main text. The vehicle described in this section was used to conduct all 9361 simulations, which were run sequentially using an automation procedure described next.

C.3 Simulation Algorithms

Algorithm 3 describes the procedure to simulate the brakeless reference, and Algorithm 4 describes the procedure to sequentially simulate the truck platoons. Both were implemented in IPG TruckMaker's ScriptControl, which uses the Tcl/Tk scripting language. Brakeless reference simulations are run separately from platoon simulations. The platoon simulations were conducted in two stages. First, the entire matrix of lead vehicle simulations is conducted, and the speed traces were stored for the follow vehicle simulations. Then, all the follow vehicle simulations were conducted.

Algorithm 3 IPG Truckmaker Simulation Automation: **Brakeless Reference**.

$\mathbf{v} \leftarrow [v_1 ... v_q]$
$\mathbf{m} \leftarrow [m_1 ... m_{qq}]$

procedure BrakelessReferenceLoop($\mathbf{v}, \mathbf{m}$)
    Lead Traffic Object Visible? $\leftarrow$ **false** ▷ Drive alone (do not follow)
    Allow Braking: $\leftarrow$ **false**
    for all $v_i \in \mathbf{v}$ **do**
        for all $m_j \in \mathbf{m}$ **do**
            Run Simulation at $v_i, m_k, DRR = 1.0$
            Save $\mathbf{v(t)}, \mathbf{m(t)}, \mathbf{x(t)}, \mathbf{P_b(t)}$... as `leadRun_`$v_i$`_`$m_j$`.erg`
        end for
    end for
end procedure

Algorithm 4 IPG Truckmaker Simulation Automation: **Platoon**.

$\mathbf{v} \leftarrow [v_1 ... v_q]$
$\mathbf{m} \leftarrow [m_1 ... m_{qq}]$
$\mathbf{DRR} \leftarrow [DRR_1 ... DRR_{qqq}]$
$\mathbf{t_0} \leftarrow 0.8$ sec ▷ Platooning timegap

procedure LeaderLoop($\mathbf{v}, \mathbf{m}$)
    Lead Traffic Object Visible? $\leftarrow$ **false** ▷ Drive alone (do not follow)
    Allow Braking: $\leftarrow$ **true**
    for all $v_i \in \mathbf{v}$ **do**
        for all $m_j \in \mathbf{m}$ **do**
            Run Simulation at $v_i, m_k, DRR = 1.0$
            Save $\mathbf{v(t)}, \mathbf{m(t)}, \mathbf{x(t)}, \mathbf{P_b(t)}$... as `leadRun_`$v_i$`_`$m_j$`.erg`
            `leadRun_`v_i`_`m_j`.ascii` $\leftarrow \mathbf{t}, \mathbf{x(t)}, \mathbf{v(t)}$ ▷ Store results for follower simulation
        end for
    end for
end procedure

procedure FollowerLoop($\mathbf{v}, \mathbf{m}, \mathbf{DRR}$)
    Lead Traffic Object Visible $\leftarrow$ **true** ▷ Platoon with traffic object
    Allow Braking: $\leftarrow$ **true**
    for all $v_i \in \mathbf{v}$ **do**
        for all $m_j \in \mathbf{m}$ **do**
            for all $m_k \in \mathbf{m}$ **do**
                for all $DRR_l \in \mathbf{DRR}$ **do**
                    Lead Traffic Object $\mathbf{t}, \mathbf{x(t)}, \mathbf{v(t)} \leftarrow$ `leadRun_`v_i`_`m_j`.ascii`
                    $v(t = 0) \leftarrow 1.025 \cdot v_i$ ▷ Ensure follower catches the traffic object
                    Run Simulation at v_i, m_k, DRR_l
                    Save $\mathbf{v(t)}, \mathbf{m(t)}, \mathbf{x(t)}, \mathbf{P_b(t)}$... as `followRun_`$v_i$`_`$m_j$`_`$m_k$`_`$DRR_l$`.erg`
                end for
            end for
        end for
    end for
end procedure

Appendix D

Test Replicates

Number of Replicates	Column Labels				
Row Labels	NCAT Phase 1	ACM Phase 1	NCAT Phase 2	ACM Phase 2	Grand Total
A1	**22**	**29**	**31**	**47**	**129**
Baseline	**4**	**6**	**7**	**14**	**31**
2T	**7**	**12**	**11**	**13**	**43**
35			3		3
50	3	8	5	3	19
100	4	4	3	3	14
opt100				3	3
opt50				4	4
4T	**8**	**9**	**10**	**12**	**39**
35	2		3		5
50	3	3	5		11
100	3	6	2		11
opt100				6	6
opt50				3	3
opt75				3	3
4T in rain		**2**		**3**	**5**
100		2			2
opt100				3	3
Merge			**1**	**2**	**3**
100			1	2	3
Cut-in	**3**		**2**	**3**	**8**
50	2				2
100	1		2	3	6
A2	**22**	**21**	**28**	**47**	**118**
Baseline	**8**	**4**	**8**	**17**	**37**
NMPC baseline				**3**	**3**
2T	**3**	**7**	**7**	**7**	**24**
35			2		2
50	2	3	3		8
100	1	4	2		7
opt100				4	4
opt50				3	3
4T	**8**	**8**	**10**	**12**	**38**
35	2		3		5
50	3	2	5		10
100	3	6	2		11
opt100				6	6
opt50				3	3
opt75				3	3
4T in rain		**2**		**3**	**5**
100		2			2
opt100				3	3
Merge			**1**	**2**	**3**
100			1	2	3
Cut-in	**3**		**2**	**3**	**8**
50	2				2
100	1		2	3	6

Figure D.1: Dataset 1 test replicates, part 1.

T13	**16**	**23**	**24**	**34**	**97**
Baseline	**2**	**4**	**6**	**7**	**19**
NMPC baseline				**3**	**3**
2T	**3**	**8**	**6**	**5**	**22**
35			2		2
50	2	3	2		7
100	1	5	2		8
opt100				3	3
opt50				2	2
4T	**8**	**9**	**9**	**11**	**37**
35	2		3		5
50	3	3	4		10
100	3	6	2		11
opt100				6	6
opt50				2	2
opt75				3	3
4T in rain		**2**		**3**	**5**
100		2			2
opt100				3	3
Merge			**1**	**2**	**3**
100			1	2	3
Cut-in	**3**		**2**	**3**	**8**
50	2				2
100	1		2	3	6
T14	**22**	**31**	**31**	**41**	**125**
Baseline	**5**	**5**	**6**	**5**	**21**
NMPC baseline				**3**	**3**
2T	**6**	**15**	**11**	**13**	**45**
35			3		3
50	2	10	5	3	20
100	4	5	3	3	15
opt100				3	3
opt50				4	4
4T	**8**	**9**	**11**	**12**	**40**
35	2		3		5
50	3	3	5		11
100	3	6	3		12
opt100				6	6
opt50				3	3
opt75				3	3
4T in rain		**2**		**3**	**5**
100		2			2
opt100				3	3
Merge			**1**	**2**	**3**
100			1	2	3
Cut-in	**3**		**2**	**3**	**8**
50	2				2
100	1		2	3	6
Grand Total	**82**	**104**	**114**	**169**	**469**

Figure D.1: Dataset 1 test replicates, part 2.

Sum of GroupCount		Spacing							
Truck	Configuration	30	40	50	75	150	258	N/A	Grand Total
A1		**3**	**4**	**3**	**3**	**1**	**3**	**6**	**23**
	Baseline							6	6
	Aligned	3	4	3	3	1	3		17
A2		**3**	**4**	**3**	**4**	**1**	**3**	**7**	**25**
	Baseline							7	7
	Aligned	3	4	3	3	1	3		17
	Aligned with Cut-ins				1				1
Grand Total		**6**	**8**	**6**	**7**	**2**	**6**	**13**	**48**

Figure D.2: Dataset 2 test replicates.

Number of Replicates		Westbound		
Truck	Spacing	West	East	Grand Total
A1		**16**	**16**	**32**
	Baseline	4	4	8
	75	7	7	14
	150	5	5	10
A2		**17**	**17**	**34**
	Baseline	5	5	10
	75	7	7	14
	150	5	5	10
Grand Total		**33**	**33**	**66**

Figure D.3: Dataset 3 test replicates.

Number of Replicates			Direction		
Truck	Lead Speed Control	Follow Gap Control	East	West	Grand Total
Control			**17**	**17**	**34**
	CC		**17**	**17**	**34**
		no follower	17	17	34
T13 (leader)			**16**	**15**	**31**
	CC		**8**	**8**	**16**
		no follower	1	1	2
		PID	4	4	8
		OPT	3	3	6
	PCC		**8**	**7**	**15**
		PID	2	2	4
		OPT	6	5	11
T14 (follower)			**15**	**14**	**29**
	CC		**7**	**7**	**14**
		no follower	1	1	2
		PID	4	4	8
		OPT	2	2	4
	PCC		**8**	**7**	**15**
		PID	3	3	6
		OPT	5	4	9
Grand Total			**48**	**46**	**94**

Figure D.4: Dataset 4 test replicates.