Thesis for the degree of doctor of philosophy in Machine and Vehicle systems

FUEL-EFFICIENT DRIVING STRATEGIES

SINA TORABI

Department of Mechanics and Maritime Sciences CHALMERS UNIVERSITY OF TECHNOLOGY Göteborg, Sweden 2020 **Fuel-efficient driving strategies** SINA TORABI ISBN 978-91-7905-362-8

 \bigodot SINA TORABI, 2020

Doktorsavhandlingar vid Chalmers tekniska högskola Ny serie nr 4829 ISSN 0346-718x

Department of Mechanics and Maritime Sciences Chalmers University of Technology SE-412 96 Göteborg Sweden Telephone: +46 (0)31-772 1000

Chalmers Digitaltryck Göteborg, Sweden 2020 "Two roads diverged in a wood, and I, I took the one less traveled by, And that has made all the difference."

- Robert Frost

Fuel-efficient driving strategies

Sina Torabi Department of Mechanics and Maritime Sciences Chalmers University of Technology

Abstract

This thesis is concerned with fuel-efficient driving strategies for vehicles driving on roads with varying topography, as well as estimation of road grade and vehicle mass for vehicles utilizing such strategies. A framework referred to as speed profile optimization (SPO), is introduced for reducing the fuel or energy consumption of single vehicles (equipped with either combustion or electric engines) and platoons of several vehicles. Using the SPO-based methods, average reductions of 11.5% in fuel consumption for single trucks, 7.5 to 12.6% energy savings in electric vehicles, and 15.8 to 17.4% average fuel consumption reductions for platoons of trucks were obtained. Moreover, SPO-based methods were shown to achieve higher savings compared to the commonly used methods for fuel-efficient driving. Furthermore, it was demonstrated that the simulations are sufficiently accurate to be transferred to real trucks. In the SPO-based methods, the optimized speed profiles were generated using a genetic algorithm for which it was demonstrated, in a discretized case, that it is able to produce speed profiles whose fuel consumption is within 2% of the theoretical optimum.

A feedforward neural network (FFNN) approach, with a simple feedback mechanism, is introduced and evaluated in simulations, for simultaneous estimation of the road grade and vehicle mass. The FFNN provided road grade estimates with root mean square (RMS) error of around 0.10 to 0.14 degrees, as well as vehicle mass estimates with an average RMS error of 1%, relative to the actual value. The estimates obtained with the FFNN outperform road grade and mass estimates obtained with other approaches.

Keywords: fuel efficiency, energy efficiency, vehicle platooning, speed profile optimization, genetic algorithms, performance analysis, mass estimation, road grade estimation, neural networks.

Acknowledgements

I am grateful to the many people who have supported me throughout my studies. Above all, I am indebted to my main supervisor Prof. Mattias Wahde who gave me the opportunity to start my PhD studies five years ago, and whose never-ending guidance and support, in this journey, helped me to complete my dissertation. I will always cherish the memories of the discussions, the projects, and the countless hours of working with you, my friend.

I am also grateful to Dr. Mauro Bellone and Prof. Pitoyo Hartono for their time, knowledge, helpful advice, and practical suggestions. I would also like to thank my fellow PhD student Luca Caltagirone for our collaboration and the long hours of correcting assignments together.

I would like to acknowledge the financial support of Vinnova/FFI and the EU Baltic sea Interreg program. I am also thankful for all the help and valuable input from the SERET project partners in setting up the experiments.

I would like to thank my colleagues and friends (past and present) at VEAS for creating a great working environment. A special thanks to Sonja who has made sure that everything is working smoothly for the PhD students.

To Leon and Kate, Juliette, Magnus, Randi, and Sabine, you should know that your support and encouragement was worth more than I can express on paper. You were always there when I needed a friend, with a word of encouragement, or a listening ear.

Thank you mom and dad, Siyavash, Nasim, Maryam, and Saman, for everything.

List of included papers

This thesis consists of the following papers. References to the papers will be made using Roman numerals.

- I. Caltagirone, L., Torabi, S., and Wahde, M., Truck Platooning based on Lead Vehicle Speed Profile Optimization and Artificial Physics. In: Proceedings of the 18th International Conference on Intelligent Transportation Systems, Las Palmas, Spain, pp. 394-399, 2015.
- II. Torabi, S., and Wahde, M., Fuel Consumption Optimization of Heavy-Duty Vehicles using Genetic Algorithms. In: Proceedings of the 2017 Congress on Evolutionary Computation, San Sebastián, Spain, pp. 29-36, 2017.
- III. Torabi, S., and Wahde, M., Fuel-efficient driving strategies for heavyduty vehicles: A platooning approach based on speed profile optimization, Journal of Advanced Transportation, 12(2018), 2018.
- IV. Torabi, S., and Wahde, M., A method for performance analysis of a genetic algorithm applied to the problem of fuel consumption minimization for heavy-duty vehicles, Applied Soft Computing, 80(2019), pp. 735-741.
- V. Torabi, S., Wahde, M., and Hartono, P., Road Grade and Vehicle Mass Estimation for Heavy-duty Vehicles Using Feedforward Neural Networks, In: Proceedings of the 4th International Conference on Intelligent Transportation Engineering, Singapore, pp. 316-321, 2019.
- VI. Torabi, S., Bellone, M., and Wahde, M., Energy minimization for an electric bus using a genetic algorithm, European Transport Research Review 12(1) (2020), pp. 1-8.

Other relevant publications co-authored by Sina Torabi:

(A) Wahde, M., Bellone, M., and Torabi, S., A method for real-time dynamic fleet mission planning for autonomous mining, Autonomous Agents and Multi-Agent Systems (2019), pp. 1-27.

Technical terms used in the thesis

Follower vehicle, 13	Genes, 25								
Genetic algorithm, 15 Speed profile optimization, 9	Hidden layer, 40								
Activation function, 40	Individual, 25								
Adaptive cruise control, 2, 12, 21 Air drag reduction 18	Kalman filtering, 37								
Artificial neural networks, 35 Artificial physics, 13	Lead vehicle, 2, 12 Leader-follower, 24								
Brute force calculations, 31	Longitudinal motion, 16								
Composite Bézier curves, 21	Model predictive control, 3, 8 Model-based estimation, 37								
Data pre-processing, 39 Derivative continuity, 27 Desired speed profile, 24 Discrete search space, 31	Optimal control, 7 Optimal control framework, 3 Optimization constraints, 29								
Driving assistance systems, 7	Penalty term, 29								
Dynamic programming, 9	Performance analyses, 31								
Encoding, 27	PID controller, 20 Platooning, 2								
Feature, 40	Platooning SPO (P-SPO), 14								
Feedforward neural network, 38	Positional continuity, 27								
Fitness measure, 28	Predictive cruise control, 8								
Follower vehicle, 2									
Fuel-efficient driving strategies, 10	Random mutation hill climbing, 26 Rectified linear unit (ReLU), 40								
Generation, 29	Recursive least squares (RLS), 37								

RMHC, 26 Road profiles, 22

Safe distance, 23 Safety constraints, 23 Sensor-based estimation, 36 Spacing policy, 12 Speed profile, 2, 23 Speed profile evaluation, 25 Speed profile optimization, 3 Standard cruise control, 2 Standardized data, 40 String stability, 12 Strings of genes, 25

Training set, 40

Validation set, 40 Vehicle model, 15

Table of Contents

1	Intr	oduction	1													
	1.1	Fuel-efficient driving	2													
	1.2	Mass and road grade estimation	4													
	1.3	Research questions	5													
	1.4	Limitations and contributions	6													
	1.5	Thesis layout	6													
2	Fue	Fuel-efficient driving														
	2.1	Optimal control framework	8													
	2.2	Speed profile optimization	9													
	2.3	Applications	10													
		2.3.1 Single vehicles	10													
		2.3.2 Platoons of vehicles	12													
3	Spe	Speed profile optimization 15														
	3.1	Modeling	15													
		3.1.1 Truck model	16													
		3.1.2 Electric minibus	19													
		3.1.3 PID controller	20													
		3.1.4 Adaptive cruise control	21													
		3.1.5 Road and speed profile representation	21													
		3.1.6 Safety constraints	23													
	3.2	Optimization	24													
		3.2.1 Evaluation \ldots \ldots \ldots \ldots \ldots \ldots \ldots \ldots	24													
		3.2.2 Optimization algorithm	25													

		3.2.3	Perfo	orman	ce ana	alysis		•		•		•	•	•	•	•	•	•	•	•	•	30
4	Para	ameter	r estii	natio	n																	35
	4.1	Sensor	r-based	l estir	natior	1																36
	4.2	Model-	-based	l estin	nation	ı																37
	4.3	Estima	ation v	with a	rtifici	al ne	ural	ne	etw	ork	s .											37
	-	4.3.1	Data	sets																		38
		4.3.2	Data	pre-r	roces	sing .																39
		4.3.3	Netw	ork a	rchite	cture	and	l tr	aiı	ning	g.	•	•		•	•		•	•		•	40
5	Disc	cussion	ı																			43
0	5.1	SPO re	- esults					_								_		_	_	_	_	43
	0.1	5.1.1	Singl	e vehi	icles .																	44
		5.1.2	Plate	ons .																		47
	5.2	SPO v	vs. MP	с																		50^{-1}
	5.3	Handli	ing of	exteri	nal tra	affic .																52
	5.4	Hetero	ogeneo	us pla	toons	5																53
	5.5	Optim	izatio	n met	hod .																	55
		5.5.1	Perfo	orman	ce ana	alysis																57
	5.6	Mass a	and ro	ad gra	ade es	timat	tion															60
	5.7	ANN v	vs. otł	ier ap	proacl	hes .																61
	5.8	Effects	s of in	accura	ite est	timat	ion						•	•	•			•				62
6	Con	clusior	ns an	d fut	ure w	vork																65
	6.1	Conclu	usions																			65
	6.2	Sugges	stions	for fu	ture v	vork				•			•	•	•			•	•	•		67
7	Sun	imary	of in	clude	d par	pers																69
	7.1	Paper	Ι																			69
	7.2	Paper	II																			70
	7.3	Paper	III											•								70
	7.4	Paper	IV											•								71
	7.5	Paper	ν											•								71
	7.6	Paper	VI					•		•		•	•	•	•	•		•	•	•	•	72
Bi	Bibliography 73												73									

INCLUDED PAPERS

Chapter

Introduction

This thesis is focused on developing and applying methods and strategies for fuel-efficient driving for trucks and electric minibuses. For conventional trucks with combustion engines (hereafter referred to as conventional trucks), fuel accounts for approximately one third of the total cost of ownership and operation. Considering that hauling companies typically own many conventional vehicles that travel an average of around 130,000 km per year [24], increasing the efficiency of their vehicles even by a few per cent can translate to significant savings for those companies. For electric vehicles, which have been introduced in recent years in the transport sector, increasing efficiency is also highly relevant, considering that current electric vehicles require frequent and time-consuming recharging cycles.

Currently, road transport is responsible for 21% of the total greenhouse gas emissions in the EU [26]. Thus, in addition to the economical benefits, technologies that increase efficiency of vehicles help to achieve the environmental goals set by governments as well as organizations such as the EU. For example, in the case of the EU, the transport sector is required to reduce greenhouse gas emissions by at least 60% by 2050 with respect to the 1990 level [25]. With these strict regulations, and considering also the expected increases in both fuel prices and the volume of transported goods, vehicle manufacturers are under pressure to take appropriate actions.

The process of improving energy efficiency and fuel efficiency of vehicles can be divided into three general categories, namely (i) improving existing vehicle designs, for example the drivetrains or aerodynamic aspects, (ii) developing entirely new vehicle designs (hybrid and fully electric engines, for example), and (iii) improving the utilization of existing vehicles, for example by applying fuel-efficient driving, which is the topic of this thesis [56].

1.1 Fuel-efficient driving

Fuel-efficient driving strategies have been developed both for groups of vehicles and for single vehicles. **Platooning**, i.e. a situation in which several vehicles drive in a linear configuration with small inter-vehicle distances, was one of the earliest strategies developed for increasing the fuel efficiency. Driving in such a configuration can reduce the fuel consumption by reducing the aerodynamic drag force experienced by the vehicles. For example, in a platoon of two trucks driving at 80 km/h at a constant spacing of 10 m, the aerodynamic drag force is reduced by around 40% [46]. In addition to the increased fuel efficiency, platooning can reduce traffic congestion and improve safety by better utilization of the road.

Most of the earlier work on platooning, both in simulations and in experiments, focused only on the aerodynamic effects of platooning at close inter-vehicle distances, when driving mainly on *flat* roads; see, for example [6, 14, 27, 37]. In these works, the lead vehicle maintained a constant speed using a standard cruise control system (CC) system while the follower vehicle maintained a constant distance using an adaptive cruise control system. However, on roads with varying topography, this approach leads to excessive acceleration and braking, and consequently neutralizes the positive effect of platooning [3, 86]. Since many roads exhibit varying topography, it is necessary to consider the impact of the road topography when developing fuel-efficient driving strategies, both for platoons of vehicles and single vehicles. The **speed profile** of a vehicle, defined as the vehicle's speed as a function of its longitudinal position along the road, thus plays an important role in any non-flat topography, both for the lead vehicle and the followers. In such situations, the speed profile must be allowed to vary according to the road's topography. An illustration of speed variation based on road topography is shown in Fig. 1.1. As can be seen in the top panel of the figure, the vehicle slows down while traversing an uphill segment and speeds up while traversing a downhill segment of a road (bottom panel).

Considering that forming or joining a platoon is not always possible, fuelefficient driving strategies should be developed for single vehicles as well. Moreover, as it turns out (see Sect. 5.3), the reduced aerodynamic drag in a



Figure 1.1: Illustration of the basic idea in speed variations based on road topography. In the top panel two speed profiles are shown, namely the Standard CC speed profile (blue line) and a fuel-efficient driving strategy (orange line). The dash lines represent the speed limits in which the fuel-efficient driving strategy operates, see Chapter 3. The bottom panel shows the road profile.

platoon formation actually plays a rather minor part, contributing only about 15% of the total fuel savings. The remaining 85% stem from following the speed profile, thus providing further motivation for considering fuel-efficient driving strategies for single vehicles.

Such strategies can be achieved in various ways which will be reviewed in this thesis (see also Chapter 2). Many of the proposed methods in the literature are based on the **optimal control framework** in general, and **model predictive control** (MPC) in particular. In these methods, the speed profile is optimized, using dynamic programming, for a rather short horizon; see e.g. [39, 41, 43, 87]. The speed profile is then tracked using inputs from the MPC controller.

An alternative approach is to use the **speed profile optimization** (SPO) framework introduced in Paper I. With the SPO approach, fuel-efficient speed profiles can be generated for road segments of any desired length. As shown in Papers II and III, the SPO method results in higher fuel savings compared to MPC-based methods for trucks equipped with internal combustion engines.

The SPO method was also extended to the case of fully electric autonomous minibuses (see Paper VI), demonstrating that the method can be applied successfully to this case as well. Moreover, the SPO method resulted in competitive energy savings (in Paper VI) relative to the methods commonly used to reduce the energy consumption of electric vehicles.

1.2 Mass and road grade estimation

For the fuel-efficient driving strategies to perform at their maximum potential, accurate knowledge of the vehicle's mass and the road grade ahead is required. Estimation strategies used to date can be categorized into two general groups, sensor-based methods and model-based methods. Sensor-based methods estimate the mass and road grade using a combination of various high-quality sensors and conventional estimation and filtering techniques, such as Kalman filtering. Model-based methods, on the other hand, use the signals received from the vehicle's communication area network (CAN) in combination with a model of the longitudinal dynamics of the vehicle for estimating the parameters, also making use of Kalman filtering.

These methods, however, either cannot be used in all driving situations (during braking, for example) or will have lower performance (for example, when certain driving styles are used; see [47]). Moreover, the performance of model-based methods deteriorates when the vehicles are following fuelefficient speed profiles, which typically contain situations involving low acceleration phases, coasting downhill without applying any engine torque, etc. Also, in certain situations such as in underground mines, it is very difficult or impossible to acquire an accurate road grade using GPS sensors.

Thus, there is a need for estimation methods which can be applied in all driving situations while also achieving higher precision than conventional methods. Moreover, more accurate estimation of the mass and the road grade may improve the performance of advanced control systems for automatic gear shifting and engine torque control as well as safety functions such as stability control and anti-lock braking systems; see e.g. [23, 54, 100, 101].

An alternative approach is to use soft computing and machine learning methods for estimating parameters of vehicles and roads. In Paper V, a feedforward neural network was implemented and applied to the case of estimating the vehicle mass and the road grade. It was shown that the performance of the neural network exceeded that of conventional methods. Moreover, the neural network's estimation is available during all driving situations, including braking, coasting, and low acceleration phases.

1.3 Research questions

The aim of this thesis has been to improve the fuel efficiency of heavy-duty vehicles equipped with internal combustion engines as well as increasing the efficiency of fully electric vehicles. More specifically, this thesis has sought to answer the following research questions (RQs):

RQ1: How should a vehicle (or a platoon) be operated in order to minimize energy consumption?

As discussed in the previous section, the fuel consumption of a conventional vehicle depends on its speed. Thus, this question refers to the problem of designing a framework with which a fuel-efficient speed profile can be generated and used by the vehicles. More importantly, the process of investigating this question may also result in ways of achieving higher efficiency than the current fuel-efficient driving approaches (such as MPC-based methods). With the prevailing trend towards electrification, it is also highly relevant to study electric vehicles in the context of this research question, in which case the energy efficiency would involve battery usage rather than fuel consumption.

RQ2: To what extent can simulation results related to energy consumption be transferred to real vehicles?

Most studies on fuel-efficient driving have been carried out only in simulations, often due to the high computational demands of MPC-based approaches. Consequently, there is a lack of information on the real-world performance of these driving strategies, which often leads to uncertainty about their effectiveness and therefore hinders their adoption in vehicles.

RQ3: How can one assess whether the performance of a fuel-efficient driving strategy is close to the theoretical optimum?

Many of the proposed fuel-efficient driving strategies are based on the optimal control framework. However, these solutions are generated for a simplified, discretized problem (where, for example, a short optimization horizon is used), which may be quite different from the original problem. In cases of the kind considered here, where no such simplifications are applied, how can one assess the performance of a fuel-efficient driving strategy? Thus, this question refers to the problem of analyzing fuel-efficient driving strategies by comparing their performance to a theoretical optimum (which can be obtained under certain assumptions). **RQ4:** How can one accurately estimate the necessary signals for implementation and execution of the driving strategy?

This question primarily concerns the problem of estimating the vehicle's mass and the road grade. In fuel-efficient driving strategies, it is commonly assumed that the mass of the vehicle and the slope of the road ahead are known with high accuracy in advance. Estimates of the vehicle mass and the road grade are utilized by the vehicle's electronic control units (ECUs) to achieve higher performance. In practice, however, these parameters are also *estimated* by the ECUs. However, as briefly mentioned in the previous section, the performance of these existing estimation methods drops during fuel-efficient driving. Thus, estimation methods with higher accuracy are needed to ensure that such driving strategies are utilized to their maximum potential.

1.4 Limitations and contributions

Here, only the longitudinal dynamics of the vehicles has been considered. Moreover, in the platoon studies presented in the papers forming this thesis, all the vehicles were equipped with the same powertrain system, but with different masses (in the case of heterogeneous platoons). Furthermore, in the simulations, it has been assumed that the motion of the vehicles was not disrupted by the surrounding traffic.

The author is the main contributor to the analysis in, and the writing of, Papers II, III, IV, V, and VI, and one of the main contributors to Paper I.

1.5 Thesis layout

Chapters 2 and 3 give the necessary background on driving strategies (Chapter 2) and the SPO methods introduced and developed for improving the efficiency of the vehicles (Chapter 3). In Chapter 4, the conventional methods used for estimation of vehicle mass and road parameters are reviewed and the proposed estimation method is presented. The main results are presented and discussed in Chapter 5, and the conclusions and future work are provided in Chapter 6. Brief summaries of the papers are given in Chapter 7.

Chapter 2

Fuel-efficient driving

The continuous refinement of vehicle technologies has been improving the energy efficiency of drivetrains over the years. Another factor that strongly affects the energy consumption of a vehicle is driving behavior, thus providing another opportunity for increased efficiency. Driving behavior can be improved via **driving assistance systems** through which drivers receive recommendations regarding the suitable target speed, something that improve their driving efficiency, see for example [7, 42, 88, 20, 75]. The increasing level of autonomy in vehicles has led to development of driving strategies that actively control the longitudinal motion of vehicles, thereby increasing the potential for energy savings even further. In this chapter, such driving strategies will be reviewed, and then the concept of speed profile optimization will be introduced.

The optimal driving strategy depends strongly on the *topography* of the road. On flat roads and under a travel time constraint, as was shown in [38], driving at constant speed is the optimal solution to the problem of fuel consumption minimization. Another simple case is that in which the slope is constant. For this case, Schwarzkopf and Leipnik [81] considered the problem of decreasing the fuel consumption of passenger cars, for which they proposed an analytical solution based on Pontryagin's maximum principle. In [28], a similar approach was proposed as well for the case of heavy-duty vehicles driving on simple road segments.

Turning now to the more realistic case of roads with varying topography, the problem of energy-efficient driving has been commonly formulated as an **optimal control** problem in which the optimization is solved numerically.



Figure 2.1: Illustration of the basic idea in the model predictive control framework. At each time instant k, an optimal control problem is solved (using dynamic programming, for example) based on the predicted state of the model, which returns an optimal control sequence (future inputs) for the entire prediction horizon p. Then, only the first of the future control inputs (Applied input) is applied to the system. At the next time instant k+1, the same procedure is repeated with updated states.

For example, Lattemann *et al.* [58] proposed the **predictive cruise control** (PCC) system, which is an upgrade on the standard cruise control system. PCC allows the speed of the vehicle to vary in a narrow range, typically around ± 5 km/h, and the system works by allowing the vehicle to lower its speed while traversing uphill segments and increase its speed while traversing downhill segments.

2.1 Optimal control framework

In order to improve the performance of PCC systems, the **model predictive control** (MPC) framework has been used in designing and tracking speed trajectories. In MPC-based approaches, an optimal speed trajectory is gen-

erated through online iterative optimization, with respect to fuel (or energy) consumption, over a typically short horizon of 2 - 5 km. The generated speed trajectory is then tracked by a controller for which the instantaneous input is generated based on the predicted state of the system, at each step. In MPC-based approaches, the optimization problem is typically solved numerically using **dynamic programming** (DP) [12].

As mentioned above, MPC is a framework that relies on iterative online solutions, for which the updating frequency depends on the discretization step of the optimal control problem. As shown in Fig. 2.1, at each time step k, an optimal control problem is formulated and solved, for the entire prediction horizon p, using DP for example. This procedure will return a sequence of control inputs of which only the first is applied to the system (between steps k and k + 1). Then, this procedure is repeated for each time instant. The discretization step, for the case of fuel-efficient driving strategies for trucks, is typically in the range of 25 to 200 meters [40, 41, 43, 87], and around 20 meters for electric vehicles since they generally operate in urban environments [63].

2.2 Speed profile optimization

In this thesis, an alternative approach for fuel-efficient driving is proposed in which an optimized speed profile is generated for longer road sections, without having to carry out the optimization iteratively at every position. In this approach, which will be referred to as the **speed profile optimization** (SPO) framework, the vehicle simply follows the generated optimized speed profile using a standard controller (for example a PID controller, see Chapter 3). In the SPO framework, the optimization is carried out using methods inspired by natural phenomena, for example genetic algorithms [44] (see Chapter 3).

One of the main advantages of the SPO-based methods, compared to the MPC-based methods, is that it is not necessary to iteratively re-calculate the solution. Therefore, once the speed profile is generated over the entire horizon, no further optimization is required.

In the SPO methods, road segments of any length can be considered in principle. In this thesis, a horizon of 10 km length has been considered, motivated by the fact that the time required to generate suitable speed profiles is typically a few minutes, making it possible to generate an optimized speed profile for the *following* 10 km section while driving over the current section.

Additionally, since the optimized speed profiles from the SPO framework act as a recipe and provide a reference speed the vehicles are not required to follow the speed profiles exactly (or even with very small errors) in contrast to the MPC-based methods. Thus, simpler controllers such as a standard PID can be used in the SPO framework.

2.3 Applications

The **fuel-efficient driving strategies** mentioned in the previous section can be used both in single vehicles and in platoons. In this section, differences between the implementations of these strategies, along with their respective performances, while driving on roads with varying topography, will be reviewed.

2.3.1 Single vehicles

Fuel-efficient driving strategies (for conventional vehicles with combustion engines) and energy-efficient strategies (for electric vehicles) have been implemented and tested rather extensively in simulations and in experiments. The savings of PCC methods and MPC-based approaches typically fall in the range of 3% to 10% for conventional vehicles (relative to highway driving at constant speed using a standard CC), and around 10% to 15% for electric vehicles (relative to human drivers operating in urban environments).

The PCC system proposed in [58] reduced the fuel consumption of a truck, driving over a road segment of 25 km, by around 3%. Hellström *et al.* [38], proposed look-ahead control (LAC) as an MPC-based upgrade of the PCC approach to solve the problem of fuel consumption minimization. In the LAC method, the optimization problem was solved using DP to generated an optimal speed profile for a truck to follow. Then, LAC was tested over 120 km of highway road, in simulation, resulting in fuel savings of 3.5% on average.

For fully electric vehicles, in contrast to the extensive work carried out for hybrid and conventional vehicles, rather few methods have been proposed for optimization of a vehicle speed profile, see for example [20, 63, 76, 77, 16, 83]. The strategies developed for driving energy efficiently for fully electric vehicles, which have mainly been tested in urban environments, are typically utilized as driving assistance systems.

Dynamic programming, i.e. the most common technique used in connection with MPC-based approaches, suffers from the *curse of dimensionality*, which is the exponential growth of the computation time with the number of states. Much work has been carried out to improve the computational efficiency of DP by, for example, eliminating the states that cannot be reached due to the physical limitations of the vehicle (see e.g. [38]). Moreover, as Henzler *et al.* [43] showed, reducing the problem of fuel-efficient driving to a convex MPC formulation improved both the computational efficiency of the method as well as the savings achieved.

In part motivated by a desire to escape the drawbacks of DP (and MPC), the concept of SPO was introduced in Paper I and tested (in simulations) for generating optimized speed profiles for the lead vehicle of a platoon. The fuel consumption of the lead truck was reduced by 15% on average, relative to standard CC, over 10 km road segments of highway roads with varying topography. In Paper II, the SPO method was implemented and tested, for a single truck, both in simulations and experiments. The analysis from Paper II showed that the results obtained in simulations are transferable to real trucks (RQ2). In other words, the fuel savings obtained in real trucks are similar to those obtained in simulations.

Moreover, the fuel savings obtained with SPO in Papers I and II compared favorably to the savings obtained by the methods mentioned earlier in this chapter. For example, in Paper II, a close comparison was made between SPO and a PCC system, in real trucks. Here, SPO improved the fuel savings by 3 percentage points, almost doubling the savings. In this comparative study, both the road and the speed range used were the same for the two methods (SPO and PCC). Furthermore, in more relaxed settings where the truck's speed was allowed to vary between 60 and 90 km/h, fuel savings of around 10% (on average, relative to standard CC) were obtained by SPO when driving on road profiles of 10 km length.

Paper VI extended the SPO method to the case of a fully electric autonomous minibus with regenerative braking. The energy consumption of the vehicle was reduced by 8.5% and 11.5% for short bus routes with slope variation, relative to the baseline case, see Fig. 5.3 in Subsect. 5.1.1. Additionally, in Paper VI, it was demonstrated that SPO enabled the minibus to carry out four more round trips on a single battery charge relative to the baseline case. The energy savings reported in Paper VI are within the range commonly found in the literature. However, the SPO results are not directly comparable to those of MPC-based approaches since the baseline case used in Paper VI is a more stringent method, offering a stronger challenge: In Paper VI, the results were compared to a (rather energy-efficient) baseline case involving autonomous operation with PID control. In other works (see e.g. [20]), a human driving cycle is used as a baseline case, which in all like-lihood would not be as efficient as an autonomous operation with a PID controller; see e.g. [65, 74].

A common critique of genetic algorithms (and similar population-based stochastic optimization methods) is their inability to guarantee that the global optimal solution will be found in finite time (see e.g. [30, 72]). In Paper IV, this issue was addressed for the problem of fuel consumption optimization of trucks by introducing a general method for *assessment* of the performance of such algorithms in slightly simplified cases where the global optimum of the problem can be determined. The results of this investigation showed that the genetic algorithm used in SPO is able to find near-optimal speed profiles for the cases considered.

2.3.2 Platoons of vehicles

The focus of most of the early work on platooning was on the concept of **string stability**, i.e. the ability of the controlled vehicles in the platoon to attenuate disturbances. Therefore, at the time, various control strategies and **spacing policies** were proposed to ensure the string stability of the platoon [73, 85, 96, 97]. The relative success of these works, centered on the practical aspects of platooning, spurned an interest in other properties such as fuel efficiency and safety.

Fuel-efficient platooning strategies are mainly concerned with maintaining a close inter-vehicle distance between vehicles to benefit from the reduced air drag. The ACC function allows the following vehicles within a platoon to follow a desired spacing policy. Thus, the combination of standard CC, for the **lead vehicle** to maintain a constant speed, and **adaptive cruise control** (ACC) for the other vehicles to keep their distance to the vehicle in front has been extensively used in platooning.

On flat highways, the CC+ACC combination for vehicle platooning showed fuel savings of up to 10% [13, 98]. However, as was shown in [1], the positive impact of the reduced air drag from driving at close inter-vehicle distances can be neutralized by slope variation, thus making the CC+ACC approach unsuitable in such situations.

There have only been few studies in which the topography of the road and its impact on fuel consumption of the entire platoon is considered. Alam *et al.* in [4] proposed a new platooning approach based on the LAC method (see Subsect. 2.3.1) which was tested on synthetic road profiles of simple uphill and downhill segments. In this approach, the LAC method is first used to generate a fuel-optimal speed trajectory for a given road segment for every vehicle in the platoon. Then, the speed trajectory that requires the largest adjustment, compared to driving at constant speed, is selected for the entire platoon to follow. This approach resulted in fuel savings of up to 14% on the downhill segments and 0.7% on the uphill segments relative to CC+ACC. A similar approach was used in [68] to generate common speed trajectories for all vehicles in a platoon. This approach reduced the fuel consumption of a platoon of four light (around 3500 kg) vehicles, driving on highways, by around 6%.

As discussed in the previous section, the SPO method for fuel-efficient driving has several advantages over MPC-based methods. In Paper I, the SPO method was used to generate optimized speed profile for the lead vehicle of a platoon. The follower vehicles used various formation control strategies adopted from the framework of **artificial physics** (AP), such as a non-linear spring damper model, modified artificial gravity, etc., to keep a safe distance (see Paper I). This approach, i.e. SPO+AP, was tested on road profiles of 10 km length (from a Swedish highway; for a more detailed discussion on road profile length, see Sect. 5.2) and reduced the fuel consumption of the entire platoon by 15% on average relative to CC+ACC. The fuel savings obtained by SPO-based platooning methods compares favorably to those obtained with MPC-based approaches. However, the utilization of the AP framework to control the **follower vehicles** of a platoon had no significant impact on the fuel reduction compared to the standard ACC. Therefore, the entire improvement in fuel savings was a result of speed profile optimizations for the lead vehicle.

Returning to MPC-based approaches, in recent work, Turri *et al.* [87] proposed a new platooning strategy where the characteristics of all vehicles were considered in the speed trajectory optimization. In this method, a common feasible speed trajectory, which can be followed by all the vehicles in the platoon, was generated using DP and tracked by the lead vehicle while the other vehicles control their distance to the vehicle in front of them. This approach, which is referred to as cooperative look-ahead control (CLAC), was tested

in simulations both for homogeneous and heterogeneous platoons of trucks on Swedish highways. The CLAC method's performance was compared to the case in which each vehicle used CC, resulting in fuel savings of 5.4% to 10.8% on average (depending on the vehicle mass and their position in the platoon).

Moreover, in order to improve the performance of MPC-based platooning approaches, in terms of both computational efficiency and fuel savings, twolayer or multi-layer hierarchical controllers have been proposed and tested in simulations. In these approaches, the first layer is responsible for generating speed trajectories for trucks, either via online optimization (i.e. during driving) [32, 62] or offline optimization [22, 33] on a longer horizon and with shorter update frequency. The speed planning results are then used in the lower layer of the controller for tracking and scheduling gear shifts. The fuel savings of these methods falls typically in the range of 4% to 15% depending on the platoon configuration and the road profile.

In Paper III, the SPO framework was extended to allow individual speed profiles for all vehicles in the platoon. In this approach, each vehicle receives and follows its own *individually* optimized speed profile *independently* of the other vehicles. In other words, in this approach which is referred to as **platooning SPO** (P-SPO), vehicles are not required to follow any spacing policies. In Paper III, the P-SPO approach was tested in simulations on road profiles of 10 km highway in Sweden, resulting in fuel savings of 15.8% (on average) for a homogeneous platoon and 16.7 – 17.4% for heterogeneous platoons (with different masses) relative to the case of CC+ACC. Moreover, in Paper III, it was shown that the fuel savings compared favorably to those of MPC-based methods in similar situations.

Chapter 3

Speed profile optimization

The main idea behind the fuel-efficient driving strategies proposed in this thesis is the offline (before driving) optimization of a vehicle's speed profile, using **genetic algorithms** [44]. With this approach, the entire speed profile is available *a priori*, before the vehicle starts driving¹. This is in contrast to most of the approaches discussed in Chapter 2, in which the optimized speed profiles are generated during driving (i.e. online) mainly using DP.

In this chapter, the methods and tools required for investigating RQs 1 and 3 are described in detail. This chapter starts with a description of the longitudinal dynamics models for conventional trucks (Papers I-IV) and for electric vehicles (Paper VI). Next, the PID controller and the ACC controller used in the simulations are presented. Then, the speed and road profile representation along with the safety constraints considered during the simulations are described. The chapter ends with a description of the method for evaluating and optimizing speed profiles, as well as the procedure used for assessing the performance of the speed profile optimization method.

3.1 Modeling

The **vehicle model** and the optimization method were implemented in a dedicated simulation environment written (by the author) in the C#.NET programming language. This simulation environment can be used both for

¹The problem of handling the presence of other vehicles is discussed in Sect. 5.3 and in Papers II, III, and VI.

single vehicles (either conventional or electric) and for a platoon of vehicles. In this thesis, only the **longitudinal motion** of a vehicle is considered both in the case of single vehicles and for platoons of vehicles.

3.1.1 Truck model

The longitudinal dynamics of a conventional truck can be expressed in the following form [3]

$$m(G)\dot{v} = F_{\rm e} - F_{\rm b} - F_{\rm d} - F_{\rm r} - F_{\rm g},$$
 (3.1)

where the expression on the left-hand side represent the vehicle's acceleration and the total inertial mass $m_{\rm G}$. The terms on the right-hand side represent the forces experienced by the vehicle, namely the engine force $F_{\rm e}$, the braking force $F_{\rm b}$, the air drag resistance force $F_{\rm d}$, the rolling resistance force $F_{\rm r}$, and the gravity force $F_{\rm g}$. The total inertial mass of a truck is computed in the following way

$$m(\mathbf{G}) = m + \frac{J_{\mathbf{w}} + \gamma_{\mathbf{G}}^2 \gamma_{\mathbf{f}}^2 \eta_{\mathbf{G}} \eta_{\mathbf{f}} J_{\mathbf{e}}}{r_{\mathbf{w}}^2}, \qquad (3.2)$$

where G is the active gear, m is the mass of the truck, $J_{\rm w}$ and $J_{\rm e}$ are the wheel and the engine inertia, respectively, $\gamma_{\rm G}$ and $\gamma_{\rm f}$ represent the gearbox and finaldrive ratios, $\eta_{\rm G}$ and $\eta_{\rm f}$ denote the gearbox and final-drive efficiencies, and $r_{\rm w}$ is the wheel radius. The various forces acting on the truck are described below.

Engine force

The engine force (F_e) acting on a truck comes from the generated torque from the truck's engine acting on the wheels, which is transferred through the drive-line, namely clutch, gearbox, etc. The engine force is computed in the following way

$$F_{\rm e} = \frac{\gamma_{\rm G} \gamma_{\rm f} \eta_{\rm G} \eta_{\rm f}}{r_{\rm w}} T_{\rm e} \equiv k_{\rm e} T_{\rm e}, \qquad (3.3)$$

where $T_{\rm e}$ is the generated torque, $\gamma_{\rm G}$ and $\gamma_{\rm f}$ are the gearbox and final-drive ratios, $\eta_{\rm G}$ and $\eta_{\rm f}$ denote the gearbox and final-drive efficiencies, and $r_{\rm w}$ is the wheel radius. The engine torque, $T_{\rm e}$, is calculated by considering the inverse dynamics of the truck: The requested acceleration $a_{\rm R}$ is computed from the speed profile, and then the required engine force is calculated by considering all the external forces acting on the truck $S = F_{\rm d} + F_{\rm r} + F_{\rm g}$ as well as the requested acceleration. Thus, by rearranging the terms in Eq. (3.1), the requested engine torque $T_{\rm e}^{\rm R}$ can be defined as

$$T_{\rm e}^{\rm R} = \frac{m({\rm G})a_{\rm R} + {\rm S}}{k_{\rm e}},\qquad(3.4)$$

The effective acceleration generated by the engine, $a_{\rm E}$, is calculated as

$$a_{\rm E} = \begin{cases} a_{\rm R} & \text{if } T_{\rm e}^{\rm R} < T_{\rm e}^{\rm max} \\ \frac{k_{\rm e} T_{\rm e}^{\rm max} - S}{m({\rm G})} & \text{if } T_{\rm e}^{\rm R} \ge T_{\rm e}^{\rm max} \end{cases}$$
(3.5)

where k_e is the torque coefficient, and T_e^{\max} is the maximum torque that can be generated by the engine. The instantaneous fuel consumption is then determined by interpolation of the torque-RPM-fuel map for the modeled vehicle, using the provided engine torque and the engine's speed.

Braking force

The braking systems of modern trucks consist of several components such as foundation (friction) brakes, engine brakes, and retarders. Given the variability of the braking systems and their controller logic, the braking torque, and consequently the braking force, is typically difficult to model in detail. Therefore, in this thesis as in many other works in the field [50, 43, 87, 67], the braking system is not modeled and, instead, it is simply assumed that the brakes can provide deceleration of down to a limit of -2.5 m/s² [80].

Air drag resistance:

The aerodynamic resistance force experienced by a single vehicle is described as

$$F_{\rm a} = \frac{1}{2} c_{\rm D} A \rho_{\rm a} v^2, \qquad (3.6)$$

where $c_{\rm D}$ is the air drag coefficient, $\rho_{\rm a}$ is the air density, A is the frontal area of the vehicle, and v is the vehicle's speed. As was discussed in the previous chapter, driving at close inter-vehicle distances (such as in a platoon, where distances of around 10 m have been used with ACC systems) reduces the



Figure 3.1: Mapping of air drag reduction as a function of inter-vehicle distance. After Turri et al. [87].

aerodynamic force. This reduction is modeled by considering a non-linear air drag ratio, $\Phi(d)$

$$F_{\rm d} = \frac{1}{2} c_{\rm D} A \rho_{\rm a} \Phi({\rm d}) v^2, \qquad (3.7)$$

where $\Phi(d)$ is a coefficient that quantifies the **air drag reduction** as a function of the inter-vehicle distance, d, when driving behind another vehicle. $\Phi(d)$ is typically modeled based on empirical results from wind-tunnel experiments as [87, 67]

$$\Phi(d_{i,i-1}) = \left(1 - \frac{C_{D,1}}{C_{D,2} + d_{i,i-1}}\right),\tag{3.8}$$

where $d_{i,i-1}$ is the i^{th} vehicle's distance to its preceding vehicle, and $C_{D,1}$ and $C_{D,2}$ are constants, obtained through regression on the experimental data in [46]. The experimental data and the fitted curve are shown in Fig. 3.1. Note that, in this model, the air drag reduction on the *preceding* vehicle is neglected (since it is much smaller than the reduction experienced by the follower vehicle).

Rolling resistance force:

The rolling resistance force, caused by the friction between the tires and the road surface, is a resistive force that can be expressed as

$$F_{\rm r} = mgc_{\rm r}\cos\alpha,\tag{3.9}$$

where m is the vehicle's mass, g is the constant of gravitational acceleration, $c_{\rm r}$ is the rolling resistance coefficient, and α is the road slope. One should note that throughout this thesis, α is positive for uphill segments of a road and negative for downhill segments.

Gravitational force:

The gravitational force plays an important role in the longitudinal dynamics of a truck and its fuel consumption, when driving on roads with varying topography, considering trucks' large masses. The gravitational force is expressed as

$$F_{\rm g} = mg\sin\alpha,\tag{3.10}$$

where, again, m is the mass of the truck, g is the gravitational acceleration, and α is the road slope. The gravitational force can either be propulsive (during downhill driving) or resistive (during uphill driving).

3.1.2 Electric minibus

The longitudinal vehicle model, used for simulating the fully electric minibus (Paper VI), is based on a computationally efficient energy consumption model proposed in [29]. In addition to providing conversion from the required power at the wheels to battery power requirements, this model provides an estimate for the battery's state-of-charge. Moreover, the model takes into account the regenerative braking system in calculations of the battery's state-of-charge. The energy consumption of the vehicle is expressed in the following general way

$$\frac{dE}{dt} = P_{\rm tot},\tag{3.11}$$

where E is the energy consumption and P_{tot} is the total power provided by the battery, computed as

$$P_{\rm tot} = P_{\rm b} \eta_{\rm RTE}, \qquad (3.12)$$

where $P_{\rm b}$, given by

$$P_{\rm b} = P_{\rm m} + P_{\rm aux},\tag{3.13}$$

is the battery's output power, η_{RTE} denotes the battery's round trip efficiency, a value which is slightly smaller than 1. The value of η_{RTE} depends on whether the battery is in charging or discharging mode (see below). P_{m} represents the required power by the electric motor whereas P_{aux} represents the auxiliary power (a value which is constant throughout the simulation) needed to keep the vehicle's auxiliary systems running.

The battery's state-of-charge, S, is updated as follows

$$\frac{dS}{dt} = -\frac{P_{\rm b}}{C},\tag{3.14}$$

where C denotes the battery's capacity in Ws.

The required power by the electric motor, $P_{\rm m}$, is calculated in the following way, depending on whether the battery is charging or discharging,

$$P_{\rm m} = \begin{cases} P_{\rm t} \eta_{\rm r}, & \text{if } P_{\rm t} < 0\\ P_{\rm t} / \eta_{\rm t}, & \text{if } P_{\rm t} > 0 \end{cases}$$
(3.15)

where $\eta_{\rm r}$ and $\eta_{\rm t}$ are the electric motor's efficiency in regenerative (charging) mode and in traction (discharging) mode, respectively. Moreover, $P_{\rm t}$ is the electric motor's required traction power which is computed by considering the longitudinal dynamics of the vehicle in the following way

$$P_{\rm t} = v \times (m(1 + C_{\rm i})a + F_{\rm a} + F_{\rm g} + F_{\rm r}), \qquad (3.16)$$

where v and a denote the vehicle's speed and acceleration, m is the mass of the vehicle, and C_i is the correction factor for rotational inertia force. The terms F_a , F_g , and F_r represent the aerodynamic drag force, gravity force, and the rolling resistance force experienced by the vehicles. These forces are calculated as described in Eqs. (3.6), (3.9), and (3.10) in the previous subsection.

3.1.3 PID controller

In both simulations and experiments, a simple **PID controller** has been used to follow a desired speed profile, both for conventional trucks and the fully electric minibus. The control output of the PID controller is computed as

$$u(t) = K_{\rm p}e(t) + K_{\rm i} \int_0^t e(\tau)d\tau + K_{\rm d}\dot{e}(t), \qquad (3.17)$$

where $K_{\rm p}, K_{\rm i}$, and $K_{\rm d}$ are the proportional, integral, and derivative gains respectively. Moreover, e(t) is the error signal which is calculated as follows

$$e(t) = v_{\rm s}(t) - v(t),$$
 (3.18)

where v_s is the reference speed and v(t) is the vehicle's instantaneous speed. The requested acceleration, a_R , which is sent to the vehicle's powertrain, is then computed by dividing the controller's output by the vehicle mass.

3.1.4 Adaptive cruise control

Adaptive cruise control (ACC), mainly used in platooning strategies, is an upgrade on the standard CC system, and it allows a vehicle to maintain a desired distance to the preceding vehicle, while following the same speed trajectory. The requested acceleration, $a_{\rm R}$, is calculated using the following expression [19]

$$a_{\rm R} = k_1 (d_{\rm i,i-1} - d_0) + k_2 (v_{\rm i-1} - v_{\rm i}), \qquad (3.19)$$

where $d_{i,i-1} = x_{i-1} - x_i$ is the inter-vehicle distance, $x_i = x_i(t)$ represents the i^{th} vehicle's position along the road, $v_i = v_i(t)$ represents the i^{th} vehicle's speed, k_1 and k_2 are the gains, and d_0 denotes the desired inter-vehicle distance.

3.1.5 Road and speed profile representation

To evaluate the SPO method's performance, both the speed profiles and the road profiles need to be modeled in the simulation environment. Two approaches were used for modeling the profiles in the papers forming this thesis. In Paper I, simple lists of two-dimensional points were used to describe the profiles, giving the elevation and the speed values every 10 meters along the road. In the other papers, a more compact representation was considered using **composite Bézier curves**, i.e. a sequence of Bézier splines; see e.g. [11]. Such a compact representation for the profiles reduces the optimization search space significantly which can lead to improved performance of the method (see Subsect. 5.5). Here, two-dimensional cubic Bézier splines of the following general form were used,



Figure 3.2: Comparison between the fitted composite Bézier curve and the original data, over a part of the road between Göteborg and Borås. The dots represent the original data and the gray curve represents the fitted splines. The vertical lines separate individual splines in the composite Bézier curve.

$$\mathbf{x}(u) \equiv (x(u), y(u))^{T} = \mathbf{P}_{0}(1-u)^{3} + 3\mathbf{P}_{1}u(1-u)^{2} + 3\mathbf{P}_{2}u^{2}(1-u) + \mathbf{P}_{3}u^{3},$$
(3.20)

where the vectors \mathbf{P}_{j} are two-dimensional control points and u is a parameter ranging from 0 to 1.

Road profiles

Two-dimensional composite Bézier curves (as described in the previous section) are used for modeling **road profiles** since only the longitudinal motion has been considered. With this representation, the two dimensions are the elevation and the longitudinal position along the road. Thus, a road profile is expressed in the following form

$$(s,z) \equiv (s_i(u), z_i(u)), \quad i = 0, \dots, n-1$$
 (3.21)

where n is the total number of splines used to model a road section. With this formulation, it becomes possible to write the elevation as z = z(s), since for any given position s along the road, the corresponding spline index and u-value can be found.

The number of splines needed to fit a composite Bézier curve to a list of position-elevation pairs can be selected in various ways. For example, it can
be chosen in such a way that the curve passes through all the data points. This approach, however, is not so suitable for two main reasons. First, fitting a curve to all the data points can introduce the noise present in the data into the road profiles. Second, if the number of splines used in fitting a Bézier curve approaches the number of data points, the search space size, taking into account the optimization procedure (see Subsect. 3.2), will exceed that of the case in which simple point lists are used. For the data sets used in this thesis, the number of splines was selected so that each spline covers roughly 500 meters of a road segment. An example of a fitted curve is shown in Fig. 3.2.

Speed profile:

Similar to the road profiles described in the previous section, the **speed profiles** are modeled using two-dimensional composite cubic Bézier curves of the following form

$$(s,v) \equiv (s_i(u), v_i(u)), \quad i = 0, \dots, n-1$$
 (3.22)

where v is the longitudinal speed of the vehicle and s is the longitudinal position of the vehicle along the road. Thus, the speed of a vehicle can be expressed as v = v(s) (similar to the elevation of the road). With this representation, the same number of splines for s as in the road profiles must be used since the speed of a vehicle is calculated based on the vehicle's current longitudinal position along the road.

3.1.6 Safety constraints

In the proposed P-SPO approach for platooning, the inter-vehicle distances between vehicles are not controlled directly. Thus, it is necessary to have **safety constraints** and to consider them during optimization. Here, the safety of a platoon is guaranteed by preventing the inter-vehicle distance from dropping below a safe distance at any time. At each time step, during the optimization, the **safe distance** is calculated as [97]

$$d_{i, i-1}^{\text{safe}}(t) = d_0 + h(t)v_i(t)$$
(3.23)

where $d_{i,i-1}^{\text{safe}}(t)$ is the minimum allowed distance between the i^{th} and $(i-1)^{\text{th}}$ vehicles at time t, d_0 is the absolute allowed minimum distance, h(t) is the

variable time headway, and $v_i(t)$ is the i^{th} vehicle's speed at time t. The variable time headway h(t) is expressed as

$$h(t) = h_0 - c_h v_r(t) \tag{3.24}$$

where $h_0 > 0$ is the (constant) minimum time headway, $c_h > 0$ is a constant, and $v_r(t) = v_{i-1} - v_i$ is the relative speed. The value of h(t) must always be positive for safety reasons, while also being prevented from becoming too large. Large values for time headway are not desirable since they can push a platoon beyond the limit where the platoon can be considered coherent. Thus in Paper III, the variable time headway values were limited to the interval [0,1] (s), and the values of h_0 and c_0 were set to 0.1 (s) and 0.2 (s²/m) respectively, as was proposed in [97].

3.2 Optimization

Here, the problem of moving a single vehicle from a given starting point to a given destination is considered for cases where the road profile is known in *advance*, so that the motion of a vehicle can be formulated as a **desired speed profile**, $v_d(s)$, defined over the entire road profile, where v_d is the desired speed of the vehicle at any given longitudinal position *s* along the road. Thus, the problem of minimizing the fuel (or energy) consumption of a vehicle involves first finding an optimal speed profile and then driving accordingly over the road (RQ1).

For platoons, two approaches have been considered. In Paper I, the **leader-follower** approach was used, in which the lead vehicle's motion is considered as an SPO problem for a single vehicle while the other (follower) vehicles control the distance to the vehicle immediately in front of them. However, in Paper III, the P-SPO method (see Subsect. 2.3.2) was used. In this method, the problem of minimizing the fuel consumption of a platoon is reduced to finding an optimal speed profile for each vehicle while ensuring a safe distance between the vehicles (see Subsect. 3.1.6), thus eliminating the problem of (active) inter-vehicle distance control.

3.2.1 Evaluation

The models and methods described in Sect. 3.1 were implemented (by the author) in a dedicated simulation environment in C#.NET for simulating

the motion of both conventional and electric vehicles. In the simulation environment, with a speed profile available, the vehicle can be made to follow it so that its fuel (or energy) consumption can be measured.

The **speed profile evaluation** for a single vehicle proceeds as follows: At each time step, the current longitudinal position of the vehicle is used to calculate the desired speed from the speed profile either by linear interpolation (in cases where a point list is used for the road profile, see Paper I), or via the spline interpolation, as described in Subsect. 3.1.5; see also Papers II, III, and VI. Thus, the speed profile is used as a lookup table from which the desired speed is calculated at any given longitudinal position along the road. The obtained desired speed is then fed to the vehicle's PID controller as its reference speed. When the vehicle reaches the end of a road profile, the fuel consumption or the energy consumption of the vehicle is returned from the simulation environment.

Turning to platoons, the evaluation of the method proceeds as follows: If the P-SPO method is used, the evaluation procedure is the same (for each vehicle) as in the case of a single vehicle case, since the vehicles follow their *individual* speed profiles *independently* (see Paper III), except that the inter-vehicle distance is also monitored to ensure the safety of the platoon (see Subsects. 3.1.6 and 3.2.2). In the case of the leader-follower approach (Paper I), only the lead vehicle has an optimized speed profile defined *a priori* which is evaluated as in the case of the single vehicles. For the follower vehicles, at every time step, the distance between the current vehicle and its preceding vehicle is calculated and used to compute the desired acceleration based on the adopted model, for example ACC. Once all the vehicles of a platoon reach the end of a road profile, their average fuel consumption is returned.

3.2.2 Optimization algorithm

The optimization of the speed profiles is carried out using genetic algorithms (GAs) [31, 44, 66], with respect either to fuel consumption or energy consumption. In the optimization algorithms used here, each **individual** (a candidate solution to the problem at hand) defines N speed profiles (in case of a platoon of N vehicles, Paper III) or one profile (N = 1) in case of a single vehicle (Papers I, II, and VI). Speed values defining the speed profiles are encoded in **strings** (referred to as chromosomes, following the usual GA nomenclature) of L numbers (**genes**). The aim of the optimization is to find



Figure 3.3: An example of speed profile tweaking in the RMHC method. In the snapshot shown, taken from the early stages of an optimization run, the initial (flat) speed profile has undergone three tweakings.

an individual whose corresponding values define speed profiles for which the fuel (or energy) consumption is minimal over a given road profile.

In Paper I, a simple genetic algorithm with a single individual (that is, a population of size 1) was used. The speed profile was modeled using a simple list that specified the desired speed at a number of discrete points (every 10 m in this case). This simple genetic algorithm which can also be referred to as **random mutation hill climbing** (RMHC) proceeds as follows: First, a constant speed profile is specified, i.e. one identical to the case in which the vehicle follows the standard CC system. Then, the vehicle's fuel consumption, f_0 , is computed using the evaluation procedure described in the previous section. Since only a single individual is considered in RMHC method, f_0 also represents the minimum fuel consumption, f_{\min} , found so far. Next, the individual is tweaked by randomly selecting a point at a random location at which the speed value is changed by a fraction β (either positive or negative). Then, the speed values within the interval $[x_t - r, x_t + r]$ where ris a randomly selected range, are changed linearly. Thus, the change $\Delta v_d(x)$ in the speed profile is computed as

$$\Delta v_{\rm d}(x) = \begin{cases} \left(1 - \left|\frac{x - x_{\rm t}}{r}\right|\right) \beta v_{\rm d}(x) & \text{if } |x - x_{\rm t}| < r \\ 0 & \text{otherwise} \end{cases}$$
(3.25)

An example of such a tweaking procedure is shown in Fig. 3.3. Finally, a smoothing step is applied to the speed profile using a simple, centered moving average of length 3. The new generated individual is then evaluated by simulating the motion of the vehicle following the corresponding speed profile. If the resulting fuel consumption, f_{new} , is lower than the current minimum fuel consumption, the new individual is kept and the value of f_{min} is updated accordingly. Otherwise, the new individual is discarded and a new tweaking is applied to the previous individual as described above. Two constraints are applied to the generated speed profiles encoded in the individuals, namely (i) the instantaneous maximum speed, v_{max} , must never exceed the road's speed limit and (ii) the average speed should always be above a certain threshold \bar{v}_{min} . In Paper I, these constraints were considered as hard constraints, i.e. if any of them were violated during the optimization, the corresponding individual was discarded immediately.

In Papers II, III, and VI, the optimization was carried out using a more standard GA, with a larger population size (M), typically around 100. In these papers, the speed profiles are represented by composite Bézier curves to improve the performance of the optimization process; see Subsect. 5.5. The individuals are encoded using floating-point numbers, where each number (gene) represents the second component (i.e. the speed, v) of a spline control point $\mathbf{P}_{i,j}$ where i is the spline index and j is the control point index for the spline in question; see also Eq. (3.20). In order to make sure that the decoded individuals result in smooth speed profiles, two requirements were considered during **encoding**. First, since a speed profile consists of several splines, the **positional continuity** (C0) of the profile should be taken into account so that a decoded individual forms a speed profile that is continuous over the entire road profile. The second requirement ensures the smoothness of the generated speed profile. This is achieved by preserving the **derivative** continuity (C1) of the speed profile in the encoding process. To make sure that these requirements are met, the following conditions must hold:

$$P_{i,3} = P_{i+1,0}$$
 for $i = 0, \dots, n-1$ (3.26)



Figure 3.4: The optimization procedure used in the SPO and P-SPO methods. The decoding step generates a speed profile (one for each chromosome in the population). The speed profiles (individuals) are then evaluated one by one, using the vehicle model along with the PID controller and the road profile, returning a fitness value, namely the inverse of the fuel consumption for the individual in question.

and

$$P_{i,3} - P_{i,2} = P_{i+1,1} - P_{i+1,0}$$
 for $i = 0, \dots, n-1$ (3.27)

Therefore, the number of parameters or genes (i.e. the length of the chromosomes) will be L = N(4 + 2(n - 1)) = N(2n + 2) where n is the number of splines. The decoding results in a set of N speed profiles that are then evaluated using the procedure described in Subsect. 3.2.1. The **fitness measure** of an individual is taken as the inverse of the fuel (or energy) consumption of the N vehicles while following their speed profiles. The flowchart of the optimization process is shown in Fig. 3.4.

As in the case of the RMHC method used in Paper I, the generated speed profiles from the standard GA used in Papers II, III, and VI must fulfill certain requirements, namely that (i) the instantaneous maximum speed must never exceed the road's speed limit, (ii) the average speed should always be above a certain threshold, and (iii) the instantaneous minimum speed should be above a certain user-defined threshold to ensure that the vehicle does not affect the traffic negatively (a condition that was not considered in Paper I). Moreover, in the case of a platoon (Paper III), two additional constraints were considered to ensure the safety and cohesion of the platoon: (iv) the inter-vehicle distance was required to be larger than the safe distance defined in Eq. (3.23) and (v) the inter-vehicle distance was prevented from exceeding a threshold (here 40 m) at all time, ensuring cohesion. In contrast to the hard constraints used in Paper I, soft constraints were applied in this case: If any of the **optimization constraints** (except for the safety measure) described above was violated, the fitness value was multiplied by a penalty term smaller than 1. For instance, in a case where the instantaneous inter-vehicle distance exceeded its maximum allowed value, the **penalty term** was calculated as follows

$$p_{d_{i,i-1}}(d) = e^{-c_d \left(\frac{d_{i,i-1}}{d_{\max}} - 1\right)^2},$$
 (3.28)

where $d_{i,i-1}$ is the maximum distance between the i^{th} and the $(i-1)^{\text{th}}$ vehicle during the evaluation, c_d is a constant, and d_{\max} is the maximum allowed inter-vehicle distance. Similar penalty terms were used for the other constraints.

In GA, after evaluating all the individuals in the population, the next generation is generated as follows: First, individuals are selected using tournament selection with probability of p_{tour} (around 0.7 to 0.8) and tournament size S_{tour} (from 2 to 5). Thus, for each selection step, a tournament with S randomly selected individuals is formed and then, with probability p_{tour} , the best individual is selected. If the best individual is not selected, it is removed from the tournament and the process is repeated again with the remaining S-1 individuals, until either an individual is selected or a single individual remains in the tournament (in which case that individual is selected). The tournament selection step is carried out M times and the resulting M/2 pairs of individuals are then subjected to single-point crossover with probability $p_{\rm cross}$ (around 0.7 to 0.9). If the crossover is not carried out, (probability $1 - p_{\text{cross}}$) the selected individuals are copied as they are. Once M new individuals have been formed, they are all subjected to mutations such that each gene is mutated with probability of $p_{\text{mut}} = p_{\text{rel}}/L$, where p_{rel} is referred to as the relative mutation rate. Thus, if $p_{\rm rel} = 1$, there will be one mutation per individual on average. The mutation is carried out either as a real-number creep mutation (with probability p_{creep}) where the creep rate is set to 10% of the full range, or is carried out as a full-range mutation with probability $(1 - p_{creep})$. Finally, elitism is applied such that one copy of the best individual overwrites the first individual in the newly formed population. The process of decoding, evaluation, selection etc., is then repeated; see also Fig. 3.4. The GA is terminated once a certain number of individuals have been evaluated, which is determined based on the desired average speed and the length of the road profile. For example, for a road profile of 10 km length, the GA can evaluate just above 1000 individuals, on a standard laptop.

3.2.3 Performance analysis

For the SPO method a GA was selected as the optimization method, despite the fact that stochastic optimization methods (such as GAs) generally cannot guarantee that an optimal solution will be found in finite time [30, 72]. While there are many optimization algorithms that do guarantee optimality, applying such algorithms in this case would require the problem to be strongly simplified (e.g. linearized). With GAs, by contrast, no such problem simplification is required.

There have been plenty of numerical studies devoted to assessing the performance of GAs (see, e.g. [8, 10, 18, 21]), i.e. RQ3 in Sect. 1.3. Such studies are generally carried out on benchmark functions for which the global optimum (or optima) are known. Thus, in such cases, the GA can be run for many times which allows one to measure its performance. However, this approach is not applicable to practical applications in which the optimum is not known, and where the size of the search space (in terms of the number of possible variable settings that can be defined) is too large (or infinite if the optimization variables take continuous values), making it impossible to obtain the global optimum through a brute force search.

In order to make a brute force search possible, so as to find an optimal solution against which the results of the GA can be compared, one can discretize the optimization variables (while maintaining the full complexity of the truck model and its evaluation). Discretizing the optimization variables makes it possible for a brute force search to find the global optimum in the *discretized* version of the original problem. Thus, in a discretized *n*dimensional problem where m different discrete levels are defined for each variable, the number of possible variable settings equals

$$n_{\text{total}} = m^n, \tag{3.29}$$



Figure 3.5: An example of a road profile (top panel), and a piecewise linear speed profile used in analyzing the performance of the GA in an SPO problem (bottom panel).

Such discretization of the search space requires two conditions to be met: First of all, since n_{total} grows very quickly with n and m, the number of dimensions and levels cannot be large. Second, the objective function (fuel or energy consumption minimization here) should ideally be continuously differentiable, or at least vary sufficiently slowly so that the discretization does not result in a significantly distorted version of the true search space; see also Subsect. 5.5.1.

Assuming that these conditions *are* fulfilled, the benchmarking proceeds as follows: First, a **discrete search space** with m levels for each of the nvariables is generated. Then, the objective function value for each of the m^n possible settings should be evaluated and stored in a process here referred to as a **brute force calculation**. From this calculation, one would then also obtain the global optimum. For benchmarking an optimization algorithm, one should then carry out a large number of runs, storing the best objective function value found in each run. Once these runs have been completed, it is possible to **analyze the performance** of the optimization algorithm



Figure 3.6: The distribution of average speeds for the $n_{\text{total}} = 43,046,721$ possible piecewise linear speed profiles for the case n = 8 and m = 9. The $n_{\text{kept}} = 3,897,389$ profiles that fall within the desired speed range 80 ± 1 km/h are indicated with two dashed lines.

by comparing the two frequency distributions of the objective value (one from the brute force calculation and one from the optimization runs) by measuring how likely it is for the optimization algorithm to find a solution which is within p% of the optimum.

Turning now to the problem of speed profile optimization, in the SPO framework the speed profiles are defined using continuous variables, either as a (discrete) point list with continuous levels (Paper I) or as Bézier curves (Papers II, III, and VI). Since the control points of a Bézier curve can be varied in non-discrete manner, the number of possible speed profiles becomes infinite. Thus, as was discussed above, the search space must be discretized in order to analyze the performance of the GA. It is then possible (though very time-consuming) to make a brute force run through *all* possible speed profiles, computing the fuel consumption over each of these profiles. In Paper IV, instead of splines, piecewise linear profiles were used. These, in turn, were defined using discrete levels as illustrated in Fig. 3.5.

The piecewise linear speed profiles are modeled by defining n equidistant points over a road profile such that the first point is at the start of the road profile and the last point at the end. Next, from a set of m discrete values, speed levels are assigned independently to each of the n points. Thus, the total number of possible speed profiles (for a given road profile) equals $n_{\text{total}} = m^n$ as in Eq. (3.29) above. In the analysis considered in Paper IV, the values used were n = 8 and m = 9.

Since the average speed of a vehicle must fulfill a certain requirement (see

Subsect. 3.2.2), the piecewise linear speed profiles must have an average speed whose value falls within a narrow range Δv around the desired average speed \bar{v} . Here, the speed range was taken as 80 ± 1 km/h. Thus, for the purpose of the brute force calculation, out of n_{total} piecewise linear speed profiles only those profiles (denoted by n_{kept}) are kept whose average speed falls within the predefined range. This procedure is illustrated in Fig. 3.6 which shows the distribution of average speeds, along with two vertical lines that define the subset of speed profiles whose average speed falls in the desired range.

Once the n_{kept} profiles have been found, the brute force calculation can be carried out. Thus, using the vehicle model described in Sect. 3.1, the vehicle's traversal over the road profile in question is simulated, and the fuel (or energy) consumption is logged. At the end of these computationally intensive calculations, the fuel consumption for every possible piecewise linear speed profile (with the constraints defined above) will be available to serve as a benchmark for the GA applied to the SPO problem.



Parameter estimation

Access to accurate estimation of road grade and vehicle mass¹ has become essential in recent years due to the higher safety requirements for vehicles [23, 54, 100, 101] as well as the introduction and ongoing improvement of fuelefficient driving strategies for which accurate estimation of parameters such as the vehicle mass and the slope of the road ahead is required. Moreover, accurate estimation can improve the performance of other control systems in a vehicle, such as regenerative braking systems [48], power management systems [45], braking systems [101], etc. Furthermore, the mass of a fully loaded truck, which stays fairly constant *during* driving, can be up to 400% of an unloaded one [23]. Thus, considering that even a mild road grade can put a serious load on the vehicle, accurate estimation of the vehicle mass becomes very important for the vehicle's safety systems such as stability control [17] and anti-lock braking systems [100].

Simultaneous estimation of road grade and vehicle mass is a challenging task considering the time-varying nature of the road grade, compared to the mass which stays rather constant during driving. The methods for estimating road grade and vehicle mass are commonly divided into two general categories: Sensor-based approaches and model-based approaches. In this chapter, such estimation techniques will be reviewed, and then an alternative estimation method based on **artificial neural networks** (ANNs) will be introduced and described (RQ4 in Sect. 1.3).

¹In addition to these two fundamental parameters, estimation of other parameters such as tire friction and axle loads, can also be relevant but will not be considered here.

4.1 Sensor-based estimation

In sensor-based estimation methods, road grade and vehicle mass are estimated using dedicated sensors. Road grade can be estimated directly using GPS receivers and inertial sensors. For example, in [9, 78] a road grade estimation approach was proposed based on the ratio of the vertical speed to the horizontal speed, both obtained through a GPS receiver. GPS-based estimation methods require high precision sensors which are not commonly available in production vehicles due to high cost. Moreover, estimates obtained from these methods are not reliable (or even available) in areas with poor signal reception (such as in tunnels or underground mines).

Other sensor-based approaches for road grade estimation use information available from the accelerometer and vehicle speed sensors; considering that the measured acceleration from the accelerometer is the sum of the acceleration due to gravity and the vehicle's longitudinal acceleration (obtained by differentiating the wheel speed), the road grade can be estimated as [69, 82]

$$\theta \approx \sin^{-1}\left(\frac{a_{\text{accelerometer}} - \frac{d}{dt}(v_{\text{wheel}})}{g}\right),$$
(4.1)

where $a_{\text{accelerometer}}$ is the acceleration reading from the accelerometer, v_{wheel} is the vehicle's speed obtained from the wheel speed, and g is the constant of gravitational acceleration. Unlike GPS-based methods, these estimation techniques are not affected by the GPS signal conditions. However, they suffer from the signal drift overtime due to unknown sensor biases and oscillation problems caused by the vehicle pitch motion. Filtering techniques can be used to mitigate these problems and to improve the performance as was shown in [47, 49]. The root mean square (RMS) error of the estimated road grade using sensor-based methods is typically around 0.5 to 1 degrees.

The vehicle mass can be estimated either using a conventional parameter estimation method (if the road grade is available) [9], or directly using dedicated sensors such as inclination sensors, angle sensors, or pressure sensors [53, 92]. For example, in modern trucks that are equipped with pneumatic and electronically controlled suspension systems, the vehicle mass can be estimated directly using the air pressure, see e.g. [53]. The RMS error of the estimated mass is typically within 2 to 5% of the actual value. These methods, however, require high-quality sensors which are not always available in production vehicles. Thus, there has been a growing interest in developing model-based estimation approaches.

4.2 Model-based estimation

Model-based methods used for estimating road grade and vehicle mass rely on the longitudinal dynamics model of a vehicle, along with the signals available from the on-board sensors via the vehicle's communication area network (CAN) such as the vehicle's speed and acceleration, engine torque, etc. Most model-based approaches use filtering techniques such as Kalman filtering methods, recursive least squares (RLS) methods, or a combination of both.

Model-based methods are commonly used for simultaneous estimation of the mass and the road grade [64, 84, 89, 94]. For example, in [89], an estimator based on RLS with multiple forgetting factors was introduced whereas in [84], a combination of the RLS algorithm and an extended Kalman filter was used for estimating the parameters. In addition, these methods can be used for single parameter estimation [79, 51]. For example, in [79] a road grade estimation method based on Kalman filtering was introduced in which the vehicle mass was assumed to be known known in advance.

The RMS error of the model-based approaches for road grade estimation is typically within 0.2 to 0.8 degrees while the estimated mass is within 2 to 7% of its actual value. These methods, however, require persistent excitation of the input signals causing their performance to suffer depending on the driving conditions and driver behavior. There have been some efforts to guarantee the persistent excitation of the signals by, for example, actively controlling the engine torque [94], or by including an anti-windup mechanism in the estimator [5]. However, in many situations, a vehicle can coast downhill without applying any engine torque, or drive within a narrow small range of accelerations. In these situations, which occur quite often when using a fuel-efficient driving strategy, the performance of the model-based methods is generally negatively affected. Moreover, there are other driving situations, such as during braking, in which the model-based approaches cannot provide any estimations due to lack of a detailed model.

4.3 Estimation with artificial neural networks

As was mentioned in the previous sections, the methods considered so far have several drawbacks. An alternative approach is to use a neural network for estimating the road grade and vehicle mass. Even though such



Figure 4.1: Road grade of the six generated road profiles used for generating the training data set for Paper V.

approaches have been extensively used in other vehicle-related parameter estimation problems for example involving the battery state-of-charge estimation in electric vehicles, see e.g. [15, 34, 36, 52, 60], very little work has been done to date on road grade and mass estimation.

In Paper V, in order to improve the performance and reliability of the estimation, a simple **feedforward neural network** (FFNN) augmented by a feedback mechanism was introduced for simultaneous estimation of the road grade and vehicle mass for a (simulated) truck driving on highways. In addition to providing estimates in all driving situations (such as during braking), the neural network achieved better performance than the model-based approaches, providing road grade estimates with RMS error of 0.10 to 0.14 degrees and mass estimates with RMS error of around 1% (relative to the actual value).

4.3.1 Data sets

The data sets used for training and evaluating neural networks were generated by simulating the motion of a truck following a speed profile over a number of road profiles, using the truck model from Subsect. 3.1.1 and the simulation procedure described in Subsect. 3.2.1. Moreover, the mass of the truck was set to different values to provide enough data points for training.

The data sets used in Paper V were generated using the following procedure: First, a set of j = 6 synthetic road segments of 10 km length were generated based on the characteristics of a typical Swedish highway. An example of the slope variation over a set of generated synthetic roads is shown



Figure 4.2: An example of a synthetic road profile (top panel) and a piecewise linear speed profile (bottom panel) generated for neural network training.

in Fig. 4.1. Then, for each of the generated road profiles, a piecewise linear speed profile (see Subsect. 3.2.3) was generated to mimic roughly a typical driving sequence by specifying a set speed every 500 meters (corresponding to a typical line-of-sight in highway driving). The set speeds were assigned such that the average speed of the vehicle equaled a certain value (80 km/h here) allowing the vehicle to vary its speed during the downhills and uphills, thus emulating a typical driving sequence. An example of a synthetic road and a piecewise linear speed profile is shown in Fig. 4.2. This procedure was repeated for all the j = 6 synthetic roads and all the k = 8 mass values, chosen from the set $\{20, 22, 24, 26, 28, 30, 32, 34\}$ metric tonnes. Thus, a total of $n = j \times k = 48$ speed profiles were generated from which the following signals were logged: (1) mass (m, metric tonnes), (2) road grade (α , degrees), (3) longitudinal speed (v, m/s), (4) longitudinal acceleration (a, m/s²), and (5) engine torque (T_e , Nm).

4.3.2 Data pre-processing

Several **data pre-processing** steps were carried out for preparing the data set for ANN training. First, the logged data from every road profile were split into 10 data elements resulting in a total of 60 (i.e. $10 \times j$) elements for each mass value. Then, these 60 data elements were sorted based on their

average road grade over the corresponding 1 km road sub-segment. Next, by going through the 60 *sorted* elements, groups of six consecutive data elements were formed. Finally, out of every group, five data elements were randomly assigned to the **training set** and the remaining one was assigned to the **validation set**. This procedure was repeated for every mass value resulting in 400 data elements in the training set and 80 elements in the validation set. This splitting procedure ensures that both the training and validation sets contain elements of different range of road grades and mass values.

Once the data sets had been formed, all the signals (both in the training and validation sets) were **standardized** using the following expression:

$$x_{i,\text{standardized}} = \frac{x_i - \bar{x_i}}{\sigma_i} \tag{4.2}$$

where x_i is the *i*th signal. \bar{x}_i and σ_i denote the average value and the standard deviation of the *i*th signal, over the *training* set. With this procedure the signals in the training set will have zero mean and unit variance. Note that the standardization of the signals for validation set makes use of the mean and standard deviation from the *training* set, in order to give a more realistic measure of the network's performance over previously unseen data.

4.3.3 Network architecture and training

In Paper V, the FFNN consisted of an input layer, two hidden layers, and an output layer. A schematic illustration of the FFNN is shown in Fig. 4.3. In total, 8 features (input signals) were fed to the FFNN. The first six features were the standardized values of v, a and T_e at both the current and previous time step (single-step lag). For the purpose of the *training*, the last two features used in the input layer were the lagged mass, m(t-1), and road grade, $\alpha(t-1)$, but with $\beta\%$ added uniform noise. However, during validation and testing, the previous mass and road grade features in the input layer were replaced by the corresponding FFNN *output* from the previous time step, denoted by $\hat{m}(t-1)$ and $\hat{\alpha}(t-1)$ in Fig. 4.3. This simple feedback mechanism allowed the FFNN to model the slowly varying nature of the road grade variable as well as the almost constant mass of the vehicle.

In the hidden layers, **rectified linear unit** (ReLU) **activation functions** were used, with the following form:



Figure 4.3: A schematic illustration of the fully connected feedforward neural network (FFNN) for estimating the vehicle's mass and the road grade used in Paper V. The input consists of the velocity, acceleration, and engine torque at the current time step, along with their lagged realizations. During the training phase, the mass and road grade at the previous time step with added uniform noise, denoted $\tilde{m}(t-1)$ and $\tilde{\alpha}(t-1)$, were fed to the FFNN. During the testing, however, estimated values from the previous step, i.e. $\hat{m}(t-1)$ and $\hat{\alpha}(t-1)$, were used. The FFNN has two hidden layers with eight and four neurons respectively, and with ReLU activation functions. The two output neurons have linear activation functions.

$$h_i = f(x_i) = \begin{cases} x_i & \text{if } x_i > 0\\ 0 & \text{if } x_i \le 0 \end{cases}$$

$$(4.3)$$

where h_i is the *i*th hidden unit's output, $x_i = \sum_{j=1}^n W_{ji}x_j$ is the weighted sum of the inputs to the same unit in the previous layer, W_{ji} is the weight between two connected units, and x_j denotes the output from units in previous layer. The activation function of the output layer was set to a linear function:

$$o_k = f(u_k) = \sum_{i=1}^m W_{ik} h_i$$
 (4.4)

where o_k is the output of the k^{th} output unit, W_{ik} is the weight between the i^{th} hidden unit in the last hidden layer and the k^{th} output unit, and h_i is the i^{th} hidden unit's output computed from Eq. (4.3).

The FFNN was trained using the Adam optimization algorithm [55] with an initial learning rate of 0.0005, a batch size of 20, and using mean-squarederror loss as the objective function with the L2 regularization method applied to every layer during the training to reduce overfitting. Moreover, the FFNN weights were initialized using the method introduced in [35].



Discussion

In this chapter, the main results are presented and discussed. In the first section, the results of using the SPO framework for reducing the fuel consumption of single (conventional) trucks and platoons of trucks, as well as reducing the energy consumption of a fully electric minibus are presented and discussed (RQs 1 and 2). Next, in Sects. 5.2, 5.3, and 5.4, various aspects of the SPO framework are discussed and compared to the methods commonly used for fuel-efficient and energy-efficient driving (see Sects. 2.1 and 2.3). Sect. 5.5 concerns the optimization methods used in SPO and investigates their performance (RQ3). Finally, in Sects. 5.6, 5.7, and 5.8, the main results of the ANN approach introduced for estimating road grade and vehicle mass are presented and discussed (RQ4).

5.1 SPO results

As mentioned in Chapter 2, SPO-based methods were applied to both conventional vehicles and electric vehicles in simulations as well as in experiments. In this section, an overview of the results obtained by SPO (related to the investigation of RQ1) is given for single vehicles (Papers II and VI) and for platoons of trucks (Papers I and III). Moreover, some of the assumptions made during the simulation and experimental studies (related to RQ2) as well as some of the practical aspects of the SPO and P-SPO methods, are discussed in some detail.



Figure 5.1: Top panel: One of the ten optimized speed profiles generated from the SPO method in Paper II. This speed profile was used when driving over the corresponding road profile (shown in the bottom panel), both in the simulations and in the experiments, shown in the bottom panel.

5.1.1 Single vehicles

Conventional trucks

In Paper II, the SPO method was used to generate fuel-efficient speed profiles for a single truck driving on 10 road profiles of 10 km length both in simulations and in on-road experiments. The fuel consumption of the truck was reduced by 11.5% (on average) in simulations and 10.2% (on average) in experiments, relative to the case in which the truck followed a constant speed profile using standard CC (see Tables I and II in Paper II). An example of a road profile and its corresponding optimized speed profile is shown in Figure 5.1.

The results obtained in Paper II showed that the fuel savings obtained from SPO are transferable to the case of real trucks despite the simplicity

5.1. SPO results

of the vehicle model used in the simulations (see RQ2). The average difference between the truck's speed in the simulations and in the experiments was around 0.6 km/h, a value which is not large enough to justify using a more complex longitudinal model in the simulations. While the differences in fuel savings obtained in the simulations and in the experiments are small in general (see Table I in Paper II), there are some exceptions. For example, an important factor explaining the differences between the simulation results and the experimental results is the (rare) inability of the truck to follow the speed profile due to the imposed limits on the requested acceleration from the truck's PID controller that were slightly different for the real truck (see Paper II) from the corresponding limits in the simulations. Furthermore, there are two other, relatively minor, effects that were not modeled in the simulation: First, the motion of the truck can sometimes be disrupted by other traffic (see also Sect. 5.3 below). This happened (briefly) over two road profiles in Paper II, namely profiles 1 and 9, and the corresponding results were therefore excluded in the computation of the averages considering that the main focus of the work in Paper II was on asserting the transferability of the results from the simulation environment to the trucks. Second, as a result of driving behind another vehicle, the aerodynamic forces acting on the truck can be somewhat decreased. While this effect is generally much smaller than the savings obtained by following the optimized speed profile (see Sect. 5.3), it nevertheless leads to some differences in the fuel savings obtained in the simulations and in the experiments.

Electric vehicles

In Paper VI, the SPO method was extended and tested (in simulations) for a fully electric minibus operating on 10 synthetic short bus routes of 2 km length. The synthetic bus routes were generated based on a bus route in Kongsberg, Norway; see Fig. 5.2. In order to evaluate the performance of the method, two case studies were considered, namely (i) a single round trip of 4 km and (ii) continuous operation of the minibus until battery depletion. The total energy consumption of the minibus was reduced by an average of 8.5% to 11.5% (see Table I in Paper VI) relative to the baseline case in which the minibus drove at a constant speed between bus stops, except for brief acceleration and deceleration phases; see Fig. 5.3 for an illustration of the baseline case. In the second case study, the number of completed round trips was increased by around 12% relative to the baseline case (see Table II in



Figure 5.2: The bus route (generated using [©] OpenStreetMap contributors ¹) considered within the Sohjoa Baltic project which is roughly 2 km long, with a total of 4 stops and with small slope variations. The synthetic bus routes used in Paper VI were generated based on this route.

Paper VI). In both case studies, the average speed of the minibus was set to be either 10 or 15 km/h in order to comply with the regulations imposed by the authorities for autonomous buses operating on the route considered within the project². However, even the higher average speed is on the lower end of the realistic average speed range for a bus operating in an urban environment [71]. Nevertheless, higher values of the average speed could be used in principle (see also Paper II or III).

In both case studies considered in Paper VI, it was assumed that the vehicle's mass is known, including the weight of the passengers entering and alighting from the minibus. Even though the mass of an empty vehicle can be estimated with high accuracy (see Paper V), for the case of a minibus, the weight of the passengers must also be considered. This can be done by combining statistical data on the typical passenger weight with on-board sensors measuring the number of passengers entering or exiting (such as the already

¹https://www.openstreetmap.org/copyright

²The study presented in Paper VI was carried out in the framework of the Sohjoa Baltic project; see www.sohjoabaltic.eu.



Figure 5.3: Baseline speed profile example, with average speed of 10 km/h (used in Paper VI) for an entire driving cycle (from start to end, and back again).

installed CCTV cameras in most public buses). A bank of speed profiles can then be generated for a set of different total masses so that the vehicle can switch between the profiles depending on its current mass. Alternatively, the speed profile (for a given mass) could be generated remotely, on demand, and then uploaded to the bus before it leaves the stop. However, such an approach would require a faster computing resource (for example, a cloudbased computing system) since generating an optimized speed profile on a standard desktop computer might take too long (up to a few minutes).

5.1.2 Platoons

As mentioned in Subsect. 2.3.2, the SPO-based methods for truck platooning can significantly reduce the fuel consumption of the entire platoon. More specifically, in Paper I, it was shown that the SPO+ACC combination reduced the fuel consumption of a homogeneous platoon of trucks by 15% on average relative to the baseline case of CC+ACC over ten road profiles of 10 km length (see Table II in Paper I). For a *heterogeneous platoon*³, however, considering that all vehicles are following the *same* speed profile using the SPO+ACC approach, fuel savings could be sub-optimal. The main idea for introducing the P-SPO method for platooning (Paper III) was indeed to allow the vehicles of a platoon to follow their own individually optimized speed profiles to achieve higher fuel savings; see also Subsect. 2.3.2.

In Paper III, both homogeneous and heterogeneous platoons of two trucks

 $^{^{3}}$ In Paper III, the heterogeneity was limited to differences in the masses of the trucks. However, heterogeneity can of course also refer to differences in the powertrains as well, see also Sect. 5.4 below.

were considered, as in many other works on platooning (see for example [1, 4, 13, 87). The P-SPO method was applied to the platoons over 10 road profiles of 10 km length, leading to fuel savings of around 15.8% relative to CC+ACC. However, compared to SPO+ACC, the fuel savings of the homogeneous platoon were improved by just 0.2 percentage point (see Table I in Paper III). Even though the expected fuel savings for a homogeneous platoon are the same using the either the SPO+ACC method or the P-SPO method, there are several advantages using the P-SPO method. First of all, no direct communication is required between the vehicles since they follow their own speed profiles which have to be uploaded to the trucks when the platoon is formed. In cases where such speed profiles are not available for a given road, the optimization must be completed before forming the platoon. In order to speed up the optimization, a database of speed profiles for common vehicle masses could be generated so that the lead vehicle's speed profile can be made available directly. Then, one could either use the SPO+ACC (which provides large fuel savings, albeit not as large as those obtained with P-SPO), or run P-SPO starting from the available speed profile. In cases where an optimized speed profile is not available for a given mass, a speed profile can be generated by interpolating between the available profiles. However, while such interpolation sometimes generates a suitable speed profile (see the top panel of Fig. 5.4), this is not always the case, as shown in the bottom panel of Fig. 5.4. Nevertheless, the interpolation can provide a good starting point for the optimizer in SPO (and P-SPO) and thus speed up the optimization.

A second advantage of the P-SPO method compared to SPO+ACC is that with the P-SPO method, the vehicles are not required to follow a particular spacing policy, which is beneficial in driving over steep uphill or downhill segments of the road, situations in which maintaining a constant distance will lead to unnecessarily large accelerations and decelerations. Moreover, in the ACC function, the desired inter-vehicle distance was set to its absolute allowed minimum value according to Eq. (3.23). However, while such a small spacing does reduce the airdrag coefficient, it might not be realistic during all phases of driving, due to potential failures in the electronic systems and sensors which would put the safety of the platoon at risk. As was shown in the Discussion section of Paper III, with a more realistic value for inter-vehicle distance (15 meters for example) in the ACC function, the total fuel savings obtained by SPO+ACC dropped by one percentage point, thus providing further motivation for instead using P-SPO.



Figure 5.4: Two examples of speed profiles generated by interpolating (for a 35-tonne truck) already optimized speed profiles (for a 30-tonne and a 40-tonne truck). Top panel: An example of successful speed profile interpolation. In this panel, the interpolated speed profile (gray curve) is similar to the optimized profile generated by SPO (black curve) for the 35-tonne truck. Bottom panel: An example of a case where the interpolated speed profile (gray curve) differs quite strongly from the optimized speed profile generated by SPO (black curve).

The improvement made by the P-SPO method compared to the SPO + ACC method is more noticeable in heterogeneous platoons, see Table II in Paper III. For such platoons, the P-SPO method reduced the fuel consumption of the platoon by 16.8 to 17.4% for different mass configurations, relative to CC+ACC. This improvement can be attributed to two factors: With the P-SPO method (i) the inter-vehicle distances are exploited more efficiently by not requiring the vehicles to follow a specific spacing policy, and (ii) the speed profiles are optimized *independently* while also considering the safety of the entire platoon.

5.2 SPO vs. MPC

As discussed in Sect. 2.3, the SPO-based methods for single vehicles, as well as the SPO+ACC and P-SPO methods for platoons, achieve higher fuel and energy savings compared to the MPC-based methods. Apart from larger savings, the SPO-based methods have other advantages over MPC: First of all, with SPO no online *iterative* re-calculation of the speed profiles is needed. Instead, the optimized speed profiles are generated for the entire stretch of road. However, regardless of the method used, the presence of other traffic (further discussed in Sect. 5.3) may interfere with the vehicles' ability to follow the speed profiles and therefore cause a platoon to lose its coherence. In such situations, if MPC-based leader-follower methods are used, the vehicles are most likely left with no speed trajectories to follow and thus, there would be no further savings for the platoon once the disturbance is gone. However, with the P-SPO method, since the vehicles have their own optimized speed profile to follow (which is available a priori), the vehicles can resume following their profiles (see Sect. 5.3).

Another advantage of using SPO is its capability of generating speed profiles for longer horizons (i.e. for longer road sections) without having to deal with high computational complexity; see below. This is in contrast to MPC-based methods where larger horizon length leads to computationally challenging optimization problems, even in offline settings [67]. One should note that, regardless of the method used for generating speed profiles, the horizon length must be selected in a way such that there is enough flexibility for the speed of the vehicle to vary, while still maintaining the required average speed.

The choice of optimization horizon length used here (10 km) is of course somewhat arbitrary; other lengths could certainly have been considered. However, that would raise the following question on algorithmic (run-time) complexity: How would the optimization method perform as a function of the horizon length? This was investigated in Paper III by randomly extracting five road profiles (from the road data considered in Papers II and III) for each of the four different profile lengths, denoted by λ , namely 5 km, 10 km, 15 km, and 20 km. Then, for each of the 5 × 4 = 20 road profiles, an optimized speed profile was generated in a very long offline run (lasting several days), so as to generate a benchmark speed profile with fuel efficiency that is likely to be near-optimal (see also Subsect. 5.5.1). Next, for each of the road profiles, a total of 30 runs were carried out with a running time **Table 5.1:** The average ratio (r_{λ}) between the fuel consumption obtained in an online run and that obtained in a long offline run, for four different road profile lengths λ (5, 10, 15, and 20 km), normalized by the value obtained for the 10 km case. The running time for the online runs were 3.5 minutes in the case of 5 km road profiles, 7 minutes for the 10 km road profiles and so on.

Road profile length (λ)	5	10	15	20
$\overline{r}_{\lambda}/\overline{r}_{10}$	0.996	1.000	0.972	0.974

Table 5.2: The fuel savings obtained with SPO with different horizon lengths, applied to the case of a single truck. The fuel savings are relative to the case of CC with a set speed of 80 km/h.

Horizon length (km)	2	5	10	20	40	50
Fuel savings $(\%)$	4.88	9.20	9.24	9.59	9.72	9.91

just below the time required to traverse the road profile in question at the desired average speed (80 km/h here). These runs are referred to as online runs. Specifically, for the 5 km road profiles, the running time was set to 3.5 minutes, whereas for the 10 km road profile it was set to 7 minutes, etc. For the $5 \times 30 = 150$ runs for each road profile length λ , the ratios $r_{\lambda} = f_{\rm on}/f_{\rm off}$ between the fuel consumption in the online runs and the offline runs were computed. Then, averages were formed over each of the 150 runs for different profile lengths, resulting in four values $\bar{r}_5, \bar{r}_{10}, \bar{r}_{15}$, and \bar{r}_{20} , which were normalized by the ratio \bar{r}_{10} (corresponding to the horizon length selected here). These values are reported in Table 5.1, indicating that the results obtained for longer profile lengths are slightly better, but that the difference is quite small, showing that the performance of the SPO method is not very sensitive to the horizon length. However the *actual* fuel savings obtained from the online runs with durations specified as above, does improve with the horizon length, as shown in Table 5.2.

In MPC-based approaches in which the vehicles are required to track the generated speed trajectories in a precise manner, the methods often control the vehicle's powertrain by calculating and planning the throttle angle and the gear shifting. As mentioned in Subsect. 2.3.2, the computational efficiency and fuel savings of MPC-based approaches were improved by adopting multi-layer hierarchical control architectures. In these approaches, a speed

trajectory is generated in the first layer via either online optimization or offline optimization. Then, in the lower layers, the problem of tracking the speed trajectory is handled. Even though the multi-layer control architecture used in recent MPC-based methods improved their performances, it added another layer of complexity to the optimization problem. By contrast, in SPO, the vehicles are not required to track the speed profiles as precisely since the speed profiles simply provide a reference speed for the vehicle to follow. Thus, there is no need for using and implementing more advanced control architectures in case of SPO.

Another advantage of using SPO-based methods is that no discretization and linearization of the model is required, thus allowing the vehicles to operate in a wider speed range. In contrast, in MPC-based approaches it is required to discretize and (often) linearize the model during the optimization process, to reduce the complexity of the problem. Thus, the vehicles are forced to operate in a narrower speed range that, consequently, leads to lower savings. In recent work based on the MPC framework, where a convex formulation of the problem has been considered (see e.g. [43, 67]), larger speed ranges and, therefore, larger fuel savings were obtained compared to common MPC-based approaches. However, the reported savings are still below the savings obtained using the SPO methods; see also the Discussion sections in Paper II and III.

5.3 Handling of external traffic

One of the main assumptions made in the simulations, which is also common in the literature (see, for example, [1, 22, 33, 87]), is that the vehicles are able to follow their speed profiles without much interference from the surrounding traffic. There are, of course, situations in which the surrounding traffic can disturb the motion of the vehicle, regardless of the fuel-efficient driving strategy used. Fortunately, for the case of single trucks or platoons driving on highways, such interference is rather rare since these vehicles are among the slowest on a highway. In fact, during the experimental study carried out in connection with Paper II, it was noted that only in two instances the truck was unable to follow the speed profile due to interference from external traffic. The interruptions lasted for 140 seconds in total, over more than 100 km of driving. Thus, for roughly 97% of the time, the vehicle *was* able to follow its speed profile. Similarly, for the case of the autonomous minibuses considered in Paper VI, interference (beyond occasional pedestrians crossing the street) was minimal since the buses had their own dedicated lane.

Since the SPO-based methods do not modify the speed profile during driving (see Sect. 5.2), the vehicles can easily resume following their speed profiles once an obstacle (for example another vehicle, or a pedestrian crossing the road) is gone. Still, several runs were carried out to investigate how much the savings are affected by such small disturbances from the surrounding traffic. In these runs, for each 10 km road profile, a single disturbance was introduced which lasted for 30 seconds of the 450 seconds total duration. The disturbance was implemented by artificially placing a slower vehicle, whose speed was randomly chosen in the interval [60, 70] km/h, in front of the vehicle under consideration, thus forcing the vehicle to abandon its speed profile and slow down to match the speed of the other vehicle. Once the disturbance disappeared, the vehicle could resume following its speed profile. In these runs, the average fuel savings dropped by 2 percentage points, from around 12% to around 10%.

In the case of platooning, the external traffic can endanger the safety of the platoon in cases where another vehicle cuts in between two vehicles. To handle such cut-in situations in the P-SPO method, as soon as the distance of a vehicle to its preceding vehicle drops below the safe distance defined in Eq. (3.23), the vehicle's ACC function will be activated to ensure that the inter-vehicle distance remains above the safe limit. Once the disturbance disappears, the vehicle will resume following its *already available* speed profile still saving a significant amount of fuel or energy whether or not the platoon remains coherent. In order to investigate how much fuel a vehicle could still save, even after the disruption of the platoon, several reruns were made. In these runs, the coefficient that quantified the air drag reduction, Eq. (3.8), was set to zero to remove the air drag reduction effect of platooning. The vehicles were still able to achieve around 85% of their maximum savings by just following their speed profiles relative to the case where they were in a platoon.

5.4 Heterogeneous platoons

In most platooning work to date, in which heterogeneous platoons have been considered, the heterogeneity comes from the difference in the masses of the vehicles [1, 4, 87]. However, there are other common types of heterogeneous



Figure 5.5: Top panel: An optimized set of speed profiles for a heterogeneous platoon of trucks equipped with different powertrains. The curves show the lead vehicle's speed profile (black) and the follower vehicle's speed profile (gray). Bottom panel: The inter-vehicle distance (solid line) between the two trucks and the minimum allowed safe distance (dashed line) computed using Eq. (3.23).

platoons where the vehicles have different *powertrains*. Thus, the P-SPO method was used to generate optimized speed profiles for a heterogeneous platoon of two trucks, with identical masses but different powertrains, driving on the 10 road profiles of 10 km length (as in Paper III). The fuel savings obtained by the P-SPO method were about 18% (on average) relative to the baseline case of CC+ACC. These results thus show another advantage of the P-SPO method to the SPO+ACC. One of the optimized set of speed profiles is shown in Fig. 5.5. It is clear from this figure that the speed profiles differ much more than the speed profiles generated for the case of a heterogeneous platoon where the vehicles have different masses instead (Figs. 3 and 4 in Paper III). These cases show the importance of optimizing a speed profile for each vehicle of a platoon rather than optimizing one profile for the entire platoon.

5.5 Optimization method

The optimization of the speed profiles has been carried out using two algorithms, namely RMHC (Paper I) and a more conventional GA (Papers II, III, and VI). The fuel savings obtained with both algorithms are quite similar in the cases considered here. However, the standard GA has the advantage of being parallelizable, using GPUs or multiple CPUs (or CPU cores), which could be useful in online optimization. Nevertheless, both methods are capable of generating a sufficiently good speed profile for the upcoming road section while driving on the current one.

The optimization variables defining the speed profiles were represented either by a simple point list or by a composite Bézier curve. In Paper II, it was shown that by modeling the speed profiles using the more compact Bézier representation, the average fuel savings for a truck were increased by 4 percentage points relative to the case where the simple point list was used. For a 10 km road profile, using composite Bézier curves, the number of optimization variables was reduced from 1000 to a range of 38 to 46, thus significantly reducing the size of the search space.

In both the RMHC and GA approaches, the initial population consisted of flat speed profiles corresponding to the case of driving using CC. Then, the question arises whether or not using random initialization of the population instead (as is common in stochastic optimization) would be a better strategy. For the case of a single truck driving on five road profiles, several runs were made in which the initial population consisted of randomly generated speed profiles (within the range). The fuel savings obtained with the random initialization of the population were about 13% on average, similar to those obtained when using CC initialization of the profiles. A comparison is shown in Fig. 5.6. As can be seen from the figure, the initialization affects neither the final fuel consumption nor its convergence.

However, starting from the flat profile has a few advantages. For instance, the flat initialization ensures that there is a solution that fulfills the requirements and constraints of the optimization problem (such as speed limits and average speed at all times) which is useful in cases where for any reason, the optimization must be terminated prematurely. Furthermore, random initialization of the speed profiles in the case of platooning could hinder the optimization progress since such profiles are likely to violate the safety constraints of the platoon and cause collisions. Additionally, from the brute force calculation defined in Sect. 3.2.3 (see also in Paper IV), it was found



Figure 5.6: Performance comparison between two initialization strategies for the GA, namely, (i) initialization with flat speed profiles (black curves), and (ii) initialization with random speed profiles (gray curves), for different road profiles.

that the fuel consumption corresponding to a CC profile is rather good, close to the average of the distribution of the possible speed profiles that fulfill the speed limits and average speed requirements; see also Fig 5.7 in which the fuel consumption of the CC profiles is indicated by a red stripe.

Table 5.3: For each of the five road profiles, this table shows the probability, over 100 runs with the standard GA, to obtain a fuel consumption value at most x % from the theoretical optimum. The columns have x = 2 % (second column), x = 1 % (third column), and x = 0.5 % (fourth column).

Road profile	$p_{2.00}$	$p_{1.00}$	$p_{0.50}$
1	0.98	0.78	0.30
2	1.00	0.99	0.67
3	0.95	0.62	0.21
4	1.00	0.97	0.82
5	1.00	0.95	0.55

5.5.1 Performance analysis

In Paper IV, the approach described in Subsect. 3.2.3 was used to assess the performance of the SPO method (RQ3) in the case of a single truck driving on five road profiles. First, the brute force calculation was carried out over the eligible speed profiles, denoted by n_{kept} (see Subsect. 3.2.3), to form the benchmark against which the performance of the GA was analyzed. Then, for each of the five road profiles, 100 runs were made using the GA, resulting in five distributions of fuel consumption values. In each run, the number of evaluated individuals was set to 1000, so that they can all be evaluated comfortably within the time required to traverse a road profile (see also Sect. 3.4 in Paper IV). The resulting distributions were then compared with the distributions from the brute force calculation, the outcome of which is shown in Fig. 5.7. From this figure, it can be seen that, in general, the distributions of the fuel consumption values obtained by the GA (in orange) are rather close to the theoretical optimum (indicated by a dashed line). A more detailed comparison is given in Table 5.3, showing the probability of finding a piecewise linear speed profile (for each road profile) whose fuel consumption is within 2%, 1%, and 0.5% of the theoretical optimum, respectively. As can be seen, the GA is almost always able to find values within 2% of the optimum.

For the purpose of performance analysis, piecewise linear speed profiles were used. An alternative approach would have been to use splines with discrete positions for the *control points*. With such an arrangement, one could



Figure 5.7: Fuel consumption distributions generated for the performance analysis. Two panels are shown for each road profile: The upper panel shows the distribution of fuel consumption values over n_{kept} speed profiles (see Subsect. 3.2.3) obtained from the brute force calculation. The red vertical line shows the fuel consumption of the CC profile. The bottom panel shows the distribution (over the 100 runs) of the fuel consumption values obtained with the standard GA. The dashed lines indicate the global minimum of the fuel consumption for the road profile in question.
still carry out the brute force calculation and the GA runs described above. However, with the requirements on positional and derivative continuity of the splines (see Subsect. 3.2.2), the total number of parameters required for N splines equals 2N + 2. Thus, for a smilar search space size as was used in the case of piecewise linear profiles, one could only have defined three splines over the 10 km road profiles, a number which is much smaller than the 18-20 splines used in the papers, and it was considered to be too small.

While the performance analysis presented in Paper IV was carried out using piecewise linear speed profiles, the best found profiles (see Fig. 3.5 for an illustration) were similar to the spline profiles, in regards to the number of local optima in the speed profile as well as their approximate locations. Moreover, the average fuel savings obtained for the piecewise linear speed profiles were 10.0%, whereas the average savings obtained for spline-based profiles were 12.8%, using the same GA parameters. Therefore, while (as expected) the spline-based speed profiles give better results, the piecewise linear profiles are quite close in terms of performance.

Even though the method presented in Paper IV for analyzing the performance of a GA (or any other population-based stochastic optimization algorithm) is applicable to other optimization problems, the method has its limitations related to the degree to which the discretized version of the problem approximates the continuous one. Discretization of the variables obviously leads to some information loss regarding the possible values of the objective function. For example, if the objective function is discontinuous with many sharp jumps, then the approach introduced here would not be suitable. However, if the objective function, denoted by $f(x_1, x_2, \ldots, x_n)$, is continuously differentiable, and if an upper bound ∇f_{max} can be found for its gradient, then one can estimate an upper bound on the deviation between the maximum value for the discretized case, $f_{\rm max}^{\rm disc}$, and the maximum value of the continuous function, $f_{\rm max}$. If the grid cell size (for discretization) is equal to Δx (and if it is the same in all *n* dimensions), then $f_{\max} \leq f_{\max}^{\text{disc}} + \nabla f_{\max} \sqrt{n} \Delta x/2$. Of course, in many cases, it may not be possible to obtain an upper bound on the gradient or to ascertain that the objective function is indeed continuously differentiable. However, considering the absence of formal approaches for benchmarking GAs (in general, real-world applications), the presented method for performance analysis can be valuable as exemplified by the speed profile optimization problem considered in Paper IV.



Figure 5.8: Test data set II, involving the highway between Göteborg and Borås. The top panel shows the average performance of the 100 trained FFNNs in estimating the vehicle mass, while the bottom panel shows their average performance in estimating the road grade. In both panels, the actual value is indicated by red curves, the average estimated values by dark blue curves, and the standard deviation of the estimated values by light blue areas.

5.6 Mass and road grade estimation

The neural network architecture introduced in Subsect. 4.3.3 for vehicle mass and road grade estimation (RQ4) was trained on the data set generated using the procedure described in Subsect. 4.3.1. Considering the stochastic nature of the optimization algorithm used for FFNN training, the estimated parameters (the vehicle mass and road grade) will slightly differ between runs. Thus, the FFNN was trained 100 times to form an average estimation error and, additionally, to identify possible situations in which the estimation is less reliable. For the training set, the average root mean square (RMS) error of the road grade was 0.0455 ± 0.0382 degrees, while RMS error of the vehicle mass was 432.9 ± 463.1 kg.

In Paper V, two additional test data sets were also considered for evaluation of the 100 trained FFNNs. The first data set, referred to as *Test data* set I, was generated using the procedure described in Subsect. 4.3.1, with six 10 km road segments (different from the ones used in the training set). Two mass values were used for test data set I, namely 25 tonnes and 30 tonnes. The second test set, referred to as *Test data set II* was generated with road segments of 10 km length taken from Paper II, where the vehicle mass value was set to 30 tonnes. The average RMS errors of the road grade estimation was around 0.10 degrees for test data set I, and 0.14 degrees for test data set II, while the RMS error of the estimated vehicle mass was around 1% (relative to the vehicle's actual mass) for both test sets; see also Table I in Paper V. The estimated road grades and the vehicle mass provided by the FFNN for test data set II are shown in Fig. 5.8.

The piecewise speed profiles used when generating the training data set were chosen to mimic roughly a typical human driving strategy aiming to reduce the fuel consumption of a vehicle. To make sure that the performance of the FFNN is not biased towards a specific driving style, the trained FFNNs were applied to the case in which the vehicle maintains a constant speed using standard CC. Thus, a synthetic road segment of 10 km length was generated, with road grade within ± 0.6 degrees (or $\pm 1\%$, defined as flat, see [70]). The average RMS error of the estimated road grade and the mass were 0.023 degrees and 303.7 kg (or 1% in relative mass), respectively. These results illustrate the generalization capability of the trained FFNNs, considering that such a driving sequence was not even included in the training phase.

5.7 ANN vs. other approaches

The performance of the FFNN in estimation compares favorably with that obtained using the methods described in Sects. 4.1 and 4.2. However, a direct comparison with model-based and sensor-based approaches on the level of implementation is difficult to make, considering the unavailability of the dedicated sensors used in the methods described in Sect. 4.1, and the lack of details regarding the settings and parameters used in the model-based approaches discussed in Sect. 4.2. Nevertheless, a comparison on the level of *results* is possible, where it becomes evident that the results obtained by FFNNs exceed those obtained using the model-based and sensor-based methods; see the previous section and Chapter 4.

Another advantage of using neural network-based models is their reliability in providing estimation during driving phases (such as braking or driving through a tunnel) for which the other approaches cannot provide any estimation. For example, in model-based methods, the estimation is suspended during braking phases of driving since an accurate longitudinal



Figure 5.9: Average change in fuel savings, indicated by the dark blue line, for cases in which the vehicle follows an optimized speed profile generated for an incorrect mass (due to estimation error) relative to the case where it follows an optimized speed profile generated for the correct mass. The light blue areas indicate the standard deviation of the change in fuel savings.

braking model is typically unavailable (see e.g. [64]). Moreover, the method presented in Paper V provides accurate estimation, also, in situations where the vehicle's longitudinal acceleration is rather low (typically in the range of -0.2 to 0.2 m/s^2 (see Fig. 4 in Paper V). In these situations, which occur quite often when driving fuel efficiently, the input signals are too weak and, therefore, the model-based methods do not perform well or might even be completely turned off due to lack of sufficient excitation (see e.g. [61]).

5.8 Effects of inaccurate estimation

In the SPO-based methods, and in other fuel-efficient driving strategies, it is assumed that the vehicle parameters, such as its mass, rolling resistance coefficient, and so on, are known with high accuracy. However, in practice, these parameters are estimated, using various methods (such as the FFNN approach introduced in Paper V) which could add a level of uncertainty to the simulation results due to estimation errors. Therefore, to investigate the effect of such errors on fuel savings, several runs were made in which the vehicle followed an optimized speed profile generated with the estimated vehicle mass (using various levels of estimation error, in the range of $\hat{m} = m \pm 10\%$) during the optimization process. The estimation errors represent either underestimated or overestimated values of the vehicle mass. For each estimation error, the fuel savings as well as the change in fuel savings relative to the case in which the mass was estimated without any error, were calculated. Then, averages were formed over changes in fuel savings for each estimation error. The results are shown in Fig. 5.9 where it can be seen that underestimation of the mass slightly increases the fuel savings in most cases. However, though not visible in Fig. 5.9, in certain cases underestimated mass values can lead to significant reduction (by up to 80%) in fuel savings. Nevertheless, the average changes in fuel savings are rather small, typically less than 1 percentage point relative to the case in which the mass was estimated correctly.

Chapter 6

Conclusions and future work

In this chapter, the main conclusions of the thesis are presented along with some suggestions for future work.

6.1 Conclusions

The main conclusions of this work are as follows:

(i) The speed profile optimization (SPO) framework, for fuel-efficient or energy-efficient driving, was introduced and tested both in simulations and in experiments (RQ1). The SPO-based methods resulted in large savings, namely an average fuel reduction of 11.5% for a single conventional truck (Paper II), 7.5 to 12.6% average energy savings in electric vehicles (Paper VI), and 15.8 to 17.4% average fuel savings for (homogeneous and heterogeneous) platoons of trucks (Papers I and III), relative to their respective baseline methods.

(ii) In a study carried out both in simulations and on real roads, the performance of SPO was evaluated on a single truck driving on a typical Swedish highways. The fuel consumption of the truck was reduced by 11.5% (on average) in the simulations and by 10.2% (on average) in the experiments. The similarity of the obtained results, which is non-trivial, shows that the simulation results are transferable to real trucks (RQ2); see Paper II.

(iii) The savings obtained using the SPO-based methods exceed, by a few

percentage points, those obtained by commonly used methods for fuel-efficient driving described in Chapter 2.

(iv) In an on-road experiment with a conventional truck, the SPO method was compared with a common PCC approach for fuel-efficient driving, in which the savings obtained by SPO surpassed those of the PCC by 3 percentage points (almost doubling the savings); see Paper II. In a simulation study, the P-SPO method was compared with an MPC-based approach, and it was shown that the P-SPO method achieved higher fuel savings by around 3 percentage points; see Paper III.

(v) The SPO-based methods perform slightly better when longer horizon lengths are considered during the optimization. However, the run-time performance is not very sensitive to the choice of the horizon length; see Sect. 5.2.

(vi) As for the optimality of the SPO solutions, the performance of the GA was analyzed using a method introduced in Paper IV. It was shown that the GA, with a restriction to piecewise linear speed profiles, was able to find profiles whose corresponding fuel consumption are within 2% of the theoretical optimum almost all the time (RQ3).

(vii) A feedforward neural network (FFNN) approach for estimating vehicle mass and road grade was introduced, and tested in simulations (RQ4). The presented FFNN provided road grade estimations with an average root mean square (RMS) error of 0.10 to 0.14 degrees, as well as vehicle mass estimations with an average RMS error of 1% (relative to the actual mass of the vehicle); see Paper V.

(viii) Apart from providing more accurate estimations compared to the model-based and sensor-based approaches described in Chapter 4, the FFNN estimates are more reliable in driving situations (such as during braking phases) in which the estimation is typically suspended or unavailable with the other methods; see Sect. 5.7 and the discussion in Paper V.

6.2 Suggestions for future work

In Paper II, the simulation results from the SPO method were validated in real trucks. Similar studies would be relevant also for the platooning approach (P-SPO) and the SPO method for electric minibuses. For example, considering the simple wind tunnel experiments used for modeling and quantifying the platooning effect (Papers I and III), it would be interesting to investigate how the results would differ with a more advanced aerodynamic model (obtained via more detailed wind tunnel experiments).

On real roads and in the presence of other traffic, it is possible that the motion of a vehicle is affected (though this is typically rather rare in highway driving, as was noted during the on-road experiments in connection with Paper II, and discussed in Sect. 5.3). Even though, with the SPO-based methods, the vehicles are able to resume following their speed profiles once the external disturbance is gone, the fuel savings will be affected. Thus, it would be valuable to extend the small study described in Sect. 5.3 on the effects of external traffic on fuel savings, to include a more detailed and realistic traffic model. Such studies would help the responsible authorities and organizations to evaluate the impact of fuel-efficient driving strategies in more realistic scenarios, and thus to achieve the environmental goals set by the EU. Moreover, these studies could help city planners to ensure that the fuel-efficient driving strategies are utilized to their maximum potential.

Vehicle automation offers the potential for substantial reductions in fuel and energy consumption and emissions by, for example, controlling the longitudinal motion of a vehicle in a more efficient way. The methods considered here can easily be adapted for automated vehicles (either conventional or fully electric vehicles). However, these reductions are not assured since they are not direct consequences of automation. In fact, vehicle automation could increase vehicle usage by improving accessibility to a wider population (such as senior adults). Therefore, the overall societal impact of vehicle automation on fuel and energy consumption and emissions is rather unclear.

In the analysis carried out in [91], the estimated change in energy consumption of heavy-duty vehicles could be between -10% (in the best-case scenario) to 80% (in the worst-case scenario). In this analysis, electric vehicles were not considered, a factor that could potentially decrease the emissions by up to 70% [2]. However, the rather short travel-range on a single battery charge for electric vehicles (especially heavy-duty vehicles) is a limiting factor. Thus, developing energy-efficient driving strategies that can offer longer



Figure 6.1: Normalized distributions of the slope values in the training data (blue curve) set and test data set II (red curve) used for training and evaluating the FFNNs in Paper V.

travel-range on a single battery charge (as in Paper VI) is of great importance for accelerating the integration of electric vehicles in the transport sector.

As shown here, and in many other studies, vehicle platooning has several advantages including, but not limited to, increased fuel (or energy) efficiency. There have been several studies on the formation and coordination of platoons, e.g. [59, 90, 57, 93, 95, 99]. However, most of these studies are concerned with *how to* form a platoon without considering whether or not it is beneficial to do so. Thus, considering also that most of the fuel savings (for a vehicle) come from following an optimized speed profile (see the last paragraph of Sect. 5.3), and that different goods have different logistic requirements, it would be important to assess the benefits of forming and joining a platoon in future studies.

The FFNN approach introduced here for vehicle mass and road grade estimation achieved high performance during the tests. However, there are a few situations where the estimated road grade from the FFNN is not very accurate, for example when the magnitude of the actual road grade is large, or in cases where the vehicle's acceleration and engine torque vary fast. This slight performance drop can be explained by the distribution of the road grades in the training set, where 67% of the data points have road grade values in the range of -0.5 to 0.5 degrees; see Fig. 6.1. Therefore, to improve the performance of the FFNN in those situations, more instances of roads with steeper slopes could be included in the training set, as could other driving styles (in addition to following the semi-efficient piecewise linear speed profiles), thus providing more variability. | Chapter

Summary of included papers

This thesis consists of six papers which are concerned with the problems of (i) fuel-efficient driving using speed profile optimization for vehicles equipped with either internal combustion engines or electric motors, driving over roads with varying topography and (ii) estimating the road grade and vehicle mass for vehicles following an optimized speed profile.

7.1 Paper I

In this paper, the concept of speed profile optimization (SPO) for fuel-efficient driving was introduced for truck platooning, where the lead vehicle followed the generated optimized speed profile while the rest of the platoon followed the lead vehicle using various algorithms such as a non-linear spring-damper model and the standard adaptive cruise control (ACC). The SPO was implemented using a simple stochastic optimization algorithm referred to as the random mutation hill climbing (RMHC) method. The performance of the SPO method, together with the vehicle-following algorithms, was tested in simulations over 10 road profiles of 10 km length (note that the SPO method can be applied to road profiles of any length, see Sect. 5.2). These approaches were then compared to the baseline case of driving with a combination of standard cruise control (CC) for the lead vehicle and ACC for the other vehicles. Relative to the baseline case, average fuel savings of around 15% for the entire platoon were achieved using the combination of SPO and the vehicle-following algorithms. Moreover, the fuel savings of the lead vehicle were 15.8% (on average) compared to the case where the truck drove at constant speed using CC.

7.2 Paper II

The main purpose of Paper II was to test the performance of the SPO method, for a single truck, both in simulations and in experiments with real trucks. In this paper, it was shown that the results from the simulations are transferable to a real truck despite the simplicity of the truck model used in the simulations. Moreover, with a more direct comparison between the SPO method and the MPC-based methods (see Sect. 5.2), it was shown that the SPO method's performance exceeds that obtained with the MPC-based methods. More specifically, the SPO method obtained 11% fuel savings, compared to typical savings of 3 to 7% achieved with MPC-based methods in cases where large variations of speed were allowed (similar to the range used in SPO). Furthermore, in a more difficult setting with narrower speed range, the SPO method outperformed a standard predictive cruise control (PCC) by around 3 percentage points, over the exact same road and speed range. Finally, in this paper, a more advanced genetic algorithm was used (compared to the RMHC method used in Paper I), and the representations of the speed and road profiles were improved by using composite Bézier curves instead of simple point lists as was used in Paper I.

7.3 Paper III

In Paper III, a method for platooning based on SPO, referred to as platooning SPO (P-SPO) was introduced and evaluated, in simulations, both for homogeneous and heterogeneous platoons of two trucks. In the latter case, the heterogeneity consisted of the trucks having different masses. The proposed P-SPO method was evaluated on 10 road profiles of 10 km length resulting in average fuel savings of 15.8% for a homogeneous platoon, and between 16.8% and 17.4% (on average) for a heterogeneous platoon, relative to the case where the lead vehicle maintained a constant speed using CC and the follower vehicle used the ACC function to control its distance to the lead vehicle. Moreover, it was shown that the P-SPO method further improved the fuel savings obtained using the SPO+ACC method introduced in Paper I, by up to 1.8 percentage points. This improvement was achieved by allowing the trucks to follow their individually assigned optimized speed profiles. In the P-SPO method, the fuel-efficient speed profiles were generated together, while considering the safety of the platoon as a hard constraint during the optimization. Furthermore, when compared to an MPC-based platooning approach, the P-SPO method increased the fuel savings by around 3 percentage points when applied on identical roads. In addition to improving the fuel savings of the platooning methods, the P-SPO method eliminates the problem of controlling the inter-vehicle distance and avoids the complexity of the multi-layer control architecture used in MPC-based approaches.

7.4 Paper IV

Paper IV presented a method for evaluating the performance of stochastic optimization algorithms (in this case GAs), in situations where the global optimum is unknown or cannot be computed, which is a common situation in practical applications. The method involves discretization of the continuous search space, which makes it possible to find the global optimum by brute forcing through the entire discretized search space. The performance of an optimization algorithm can then be assessed by comparing its best solutions to the global optimum of the discretized problem. This procedure was applied to the problem of speed profile optimization of trucks using GAs, in which the optimization must be carried out under a time constraint. It was shown that the GA was able to find near-optimal solutions for the considered cases with the optimized speed profiles having objective function values within 2% of the global optimum.

7.5 Paper V

In Paper V, an artificial neural network (ANN) approach was presented for solving the problem of simultaneous estimation of road grade and vehicle mass. The ANN was applied to the case of simulated trucks driving on highways while following fuel-efficient speed profiles. It provided road grade estimates with root mean square (RMS) error of 0.10 to 0.14 degrees, and mass estimates with an average RMS error of 1% (relative to the vehicle's actual mass). The ANN was applied to two data sets, one with data obtained via simulating a truck driving on synthetic roads and one based on real roads. Moreover, the estimates obtained from the ANN outperformed those obtained with model-based approaches (such as Kalman filtering and recursive least squares methods) whose RMS errors are around 0.2 degrees for road grade estimates, and around 2% for mass. Moreover, in contrast to model-based approaches which cannot be used in certain driving situations (such as during braking), the ANN method can be used at all times.

7.6 Paper VI

In Paper VI, the SPO framework introduced in Paper I was extended to the case of energy optimization of a fully electric autonomous minibus with regenerative braking. The performance of the method was evaluated over 10 road profiles of 2 km length in two case studies, (i) optimizing the speed profile for a round trip of 4 km driving in total and (ii) continuous simulation of the electric bus until battery depletion, with mass variations representing passengers entering and alighting. In case study I, the energy consumption of the minibus was reduced by between 8.5% and 11.5% relative to the baseline case in which it followed a constant speed with short acceleration and deceleration phases. In case study II, it was shown that, using the SPO method, the minibus increased its total number of round trips by around 10% on a single battery charge compared to the baseline case.

Bibliography

- A. AL ALAM, A. GATTAMI, AND K. H. JOHANSSON, An experimental study on the fuel reduction potential of heavy duty vehicle platooning, in Intelligent Transportation Systems (ITSC), 2010 13th International IEEE Conference on, IEEE, 2010, pp. 306–311.
- [2] A. AL-SAMARI, Study of emissions and fuel economy for parallel hybrid versus conventional vehicles on real world and standard driving cycles, Alexandria Engineering Journal, 56 (2017), pp. 721–726.
- [3] A. ALAM, *Fuel-Efficient Heavy-Duty Vehicle Platooning*, PhD thesis, KTH Royal Institute of Technology, Stockholm, May 2014.
- [4] A. ALAM, J. MÅRTENSSON, AND K. H. JOHANSSON, Look-ahead cruise control for heavy duty vehicle platooning, in 16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013), Oct 2013, pp. 928–935.
- [5] S. ALTMANNSHOFER AND C. ENDISCH, Robust vehicle mass and driving resistance estimation, in 2016 American Control Conference (ACC), July 2016, pp. 6869–6874.
- [6] L. ALVAREZ AND R. HOROWITZ, Safe platooning in automated highway systems, California Partners for Advanced Transit and Highways (PATH), (1997).
- [7] C. ANDRIEU AND G. S. PIERRE, Using statistical models to characterize eco-driving style with an aggregated indicator, in 2012 IEEE Intelligent Vehicles Symposium, June 2012, pp. 63–68.

- [8] T. BÄCK AND H.-P. SCHWEFEL, An overview of evolutionary algorithms for parameter optimization, Evolutionary computation, 1 (1993), pp. 1–23.
- [9] H. S. BAE, J. RYU, AND J. C. GERDES, Road grade and vehicle parameter estimation for longitudinal control using gps, in Proceedings of the IEEE Conference on Intelligent Transportation Systems, 2001, pp. 25–29.
- [10] P. J. BALLESTER, J. STEPHENSON, J. N. CARTER, AND K. GAL-LAGHER, Real-parameter optimization performance study on the cec-2005 benchmark with spc-pnx, in 2005 IEEE Congress on Evolutionary Computation, vol. 1, IEEE, 2005, pp. 498–505.
- [11] R. H. BARTELS, J. C. BEATTY, AND B. A. BARSKY, An Introduction to Splines for Use in Computer Graphics & Geometric Modeling, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1987.
- [12] R. BELLMAN, Dynamic Programming, Princeton University Press, Princeton, NJ, USA, 1 ed., 1957.
- [13] C. BONNET AND H. FRITZ, Fuel consumption reduction in a platoon: Experimental results with two electronically coupled trucks at close spacing, tech. rep., SAE Technical Paper, 2000.
- [14] K. S. CHANG, J. KARL HEDRICK, W.-B. ZHANG, P. VARAIYA, M. TOMIZUKA, AND S. E. SHLADOVER, Automated highway system experiments in the path program, Journal of Intelligent Transportation Systems, 1 (1993), pp. 63–87.
- [15] M. CHARKHGARD AND M. FARROKHI, State-of-charge estimation for lithium-ion batteries using neural networks and ekf, IEEE transactions on industrial electronics, 57 (2010), pp. 4178–4187.
- [16] H. CHEN, L. GUO, H. DING, Y. LI, AND B. GAO, Real-time predictive cruise control for eco-driving taking into account traffic constraints, IEEE Transactions on Intelligent Transportation Systems, 20 (2019), pp. 2858–2868.

- [17] W. CHEN, D. TAN, AND L. ZHAO, Vehicle sideslip angle and road friction estimation using online gradient descent algorithm, IEEE Transactions on Vehicular Technology, 67 (2018), pp. 11475–11485.
- [18] K. A. DE JONG, Analysis of the behavior of a class of genetic adaptive systems, tech. rep., 1975.
- [19] Q. DENG AND X. MA, A fast algorithm for planning optimal platoon speeds on highway, IFAC Proceedings Volumes, 47 (2014), pp. 8073– 8078.
- [20] W. DIB, A. CHASSE, P. MOULIN, A. SCIARRETTA, AND G. CORDE, Optimal energy management for an electric vehicle in eco-driving applications, Control Engineering Practice, 29 (2014), pp. 299–307.
- [21] J. G. DIGALAKIS AND K. G. MARGARITIS, An experimental study of benchmarking functions for genetic algorithms, International Journal of Computer Mathematics, 79 (2002), pp. 403–416.
- [22] F. DONATANTONIO, A. DAMATO, I. ARSIE, AND C. PIANESE, A multi-layer control hierarchy for heavy duty vehicles with off-line dual stage dynamic programming optimization, Transportation Research Part C: Emerging Technologies, 92 (2018), pp. 486–503.
- [23] M. DRUZHININA, A. STEFANOPOULOU, AND L. MOKLEGAARD, Adaptive continuously variable compression braking control for heavyduty vehicles, Journal of dynamic systems, measurement, and control, 124 (2002), pp. 406–414.
- [24] EUROPEAN COMMISSION, Reduction and Testing of Greenhouse Gas Emissions from Heavy Duty Vehicles, The European Commission, 2011.
- [25] —, Roadmap to a single european transport area: towards a competitive and resource-efficient transport system, 2011.
- [26] —, EU transport in figures statistical pocketbook, 2016.
- [27] H. FRITZ, Longitudinal and lateral control of heavy duty trucks for automated vehicle following in mixed traffic: experimental results from the chauffeur project, in Proceedings of the 1999 IEEE International

Conference on Control Applications (Cat. No.99CH36328), vol. 2, 1999, pp. 1348–1352 vol. 2.

- [28] A. FRÖBERG, E. HELLSTRÖM, AND L. NIELSEN, Explicit fuel optimal speed profiles for heavy trucks on a set of topographic road profiles, tech. rep., SAE Technical Paper, 2006.
- [29] K. N. GENIKOMSAKIS AND G. MITRENTSIS, A computationally efficient simulation model for estimating energy consumption of electric vehicles in the context of route planning applications, Transportation Research Part D: Transport and Environment, 50 (2017), pp. 98–118.
- [30] F. W. GLOVER AND G. A. KOCHENBERGER, Handbook of metaheuristics, vol. 57, Springer Science & Business Media, 2006.
- [31] D. E. GOLDBERG, *Genetic algorithms*, Pearson Education India, 2006.
- [32] G. GUO AND Q. WANG, Fuel-efficient en route speed planning and tracking control of truck platoons, IEEE Transactions on Intelligent Transportation Systems, 20 (2019), pp. 3091–3103.
- [33] A. HAMEDNIA, N. K. SHARMA, N. MURGOVSKI, AND J. FREDRIKS-SON, Computationally efficient algorithm for eco-driving over long lookahead horizons, 2020.
- [34] M. A. HANNAN, M. S. H. LIPU, A. HUSSAIN, M. H. SAAD, AND A. AYOB, Neural network approach for estimating state of charge of lithium-ion battery using backtracking search algorithm, Ieee Access, 6 (2018), pp. 10069–10079.
- [35] K. HE, X. ZHANG, S. REN, AND J. SUN, Delving deep into rectifiers: Surpassing human-level performance on imagenet classification, in Proceedings of the IEEE international conference on computer vision, 2015, pp. 1026–1034.
- [36] W. HE, N. WILLIARD, C. CHEN, AND M. PECHT, State of charge estimation for li-ion batteries using neural network modeling and unscented kalman filter-based error cancellation, International Journal of Electrical Power & Energy Systems, 62 (2014), pp. 783–791.

- [37] J. K. HEDRICK, D. MCMAHON, V. NARENDRAN, AND D. SWA-ROOP, Longitudinal vehicle controller design for ivhs systems, in 1991 American Control Conference, June 1991, pp. 3107–3112.
- [38] E. HELLSTRÖM, J. ÅSLUND, AND L. NIELSEN, Design of an efficient algorithm for fuel-optimal look-ahead control, Control Engineering Practice, 18 (2010), pp. 1318–1327.
- [39] —, Horizon length and fuel equivalents for fuel-optimal look-ahead control, IFAC Proceedings Volumes, 43 (2010), pp. 360–365.
- [40] E. HELLSTRÖM, A. FRÖBERG, AND L. NIELSEN, A real-time fueloptimal cruise controller for heavy trucks using road topography information, tech. rep., SAE Technical Paper, 2006.
- [41] E. HELLSTRÖM, M. IVARSSON, J. ÅSLUND, AND L. NIELSEN, Lookahead control for heavy trucks to minimize trip time and fuel consumption, Control Engineering Practice, 17 (2009), pp. 245 – 254.
- [42] M. HENZLER, M. BUCHHOLZ, AND K. DIETMAYER, Online velocity trajectory planning for manual energy efficient driving of heavy duty vehicles using model predictive control, in 17th International IEEE Conference on Intelligent Transportation Systems (ITSC), Oct 2014, pp. 1814–1819.
- [43] —, Optimal parameter selection of a model predictive control algorithm for energy efficient driving of heavy duty vehicles, in 2015 IEEE Intelligent Vehicles Symposium (IV), June 2015, pp. 742–748.
- [44] J. H. HOLLAND, Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence., U Michigan Press, 1975.
- [45] Y. HUANG, H. WANG, A. KHAJEPOUR, H. HE, AND J. JI, Model predictive control power management strategies for hevs: A review, Journal of Power Sources, 341 (2017), pp. 91–106.
- [46] W. HUCHO, Aerodynamics of Road Vehicles: From Fluid Mechanics to Vehicle Engineering, Elsevier Science, 2013.

- [47] J. JAUCH, J. MASINO, T. STAIGER, AND F. GAUTERIN, Road grade estimation with vehicle-based inertial measurement unit and orientation filter, IEEE Sensors Journal, 18 (2018), pp. 781–789.
- [48] F. JI, Y. PAN, Y. ZHOU, F. DU, Q. ZHANG, AND G. LI, Energy recovery based on pedal situation for regenerative braking system of electric vehicle, Vehicle system dynamics, 58 (2020), pp. 144–173.
- [49] K. JO, J. KIM, AND M. SUNWOO, Real-time road-slope estimation based on integration of onboard sensors with gps using an immpda filter, IEEE Transactions on Intelligent Transportation Systems, 14 (2013), pp. 1718–1732.
- [50] L. JOHANNESSON, M. NILSSON, AND N. MURGOVSKI, Look-ahead vehicle energy management with traffic predictions, IFAC-PapersOnLine, 48 (2015), pp. 244–251.
- [51] D. JUNG AND G. CHOI, A new adaptive mass estimation approach of heavy truck based on engine torque local convex minimum characteristic at low speeds, Energies, 13 (2020), p. 1649.
- [52] L. KANG, X. ZHAO, AND J. MA, A new neural network model for the state-of-charge estimation in the battery degradation process, Applied Energy, 121 (2014), pp. 20–27.
- [53] R. KARLSSON, A. ANDERSSON, AND J. JOSEFSSON, *Estimation of the load of a vehicle*, December, 1 2015. US Patent 8,630,767.
- [54] D. KIM, S. B. CHOI, AND M. CHOI, Integrated vehicle mass estimation for vehicle safety control using the recursive least-squares method and adaptation laws, Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering, 229 (2015), pp. 14– 24.
- [55] D. P. KINGMA AND J. BA, Adam: A method for stochastic optimization, CoRR, abs/1412.6980 (2014).
- [56] S. KOBAYASHI, S. PLOTKIN, AND S. K. RIBEIRO, *Energy efficiency technologies for road vehicles*, Energy Efficiency, 2 (2009), pp. 125–137.

- [57] J. LARSON, C. KAMMER, K.-Y. LIANG, AND K. H. JOHANSSON, *Coordinated route optimization for heavy-duty vehicle platoons*, in 16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013), IEEE, 2013, pp. 1196–1202.
- [58] F. LATTEMANN, K. NEISS, S. TERWEN, AND T. CONNOLLY, *The predictive cruise control a system to reduce fuel consumption of heavy duty trucks*, in SAE Technical Paper, SAE International, 10 2004.
- [59] K.-Y. LIANG, J. MÅRTENSSON, AND K. H. JOHANSSON, Heavyduty vehicle platoon formation for fuel efficiency, IEEE Transactions on Intelligent Transportation Systems, 17 (2015), pp. 1051–1061.
- [60] H.-T. LIN, T.-J. LIANG, AND S.-M. CHEN, Estimation of battery state of health using probabilistic neural network, IEEE Transactions on Industrial Informatics, 9 (2012), pp. 679–685.
- [61] N. LIN, C. ZONG, AND S. SHI, The method of mass estimation considering system error in vehicle longitudinal dynamics, Energies, 12 (2019), p. 52.
- [62] J. LIU, B. PATTEL, A. S. DESAI, E. HODZEN, AND H. BORHAN, Fuel efficient control algorithms for connected and automated line-haul trucks, in 2019 IEEE Conference on Control Technology and Applications (CCTA), IEEE, 2019, pp. 730–737.
- [63] D. MAAMRIA, K. GILLET, G. COLIN, Y. CHAMAILLARD, AND C. NOUILLANT, On the use of dynamic programming in eco-driving cycle computation for electric vehicles, in 2016 IEEE Conference on Control Applications (CCA), IEEE, 2016, pp. 1288–1293.
- [64] M. L. MCINTYRE, T. J. GHOTIKAR, A. VAHIDI, X. SONG, AND D. M. DAWSON, A two-stage lyapunov-based estimator for estimation of vehicle mass and road grade, IEEE Transactions on Vehicular Technology, 58 (2009), pp. 3177–3185.
- [65] A. C. MERSKY AND C. SAMARAS, Fuel economy testing of autonomous vehicles, Transportation Research Part C: Emerging Technologies, 65 (2016), pp. 31–48.
- [66] M. MITCHELL, An introduction to genetic algorithms, MIT press, 1998.

- [67] N. MURGOVSKI, B. EGARDT, AND M. NILSSON, Cooperative energy management of automated vehicles, Control Engineering Practice, 57 (2016), pp. 84–98.
- [68] B. NEMETH AND P. GASPAR, Optimised speed profile design of a vehicle platoon considering road inclinations, IET Intelligent Transport Systems, 8 (2014), pp. 200–208.
- [69] M. NILSSON AND E. OHLUND, Estimation of road inclination, January, 14 2014. US Patent 9,200,898.
- [70] A. ODRIGO, M. EL-GINDY, P. PETTERSSON, Z. NEDĚLKOVÁ, P. LINDROTH, AND F. ÖIJER, *Design and development of a road profile generator*, International Journal of Vehicle Systems Modelling and Testing, 11 (2016), pp. 217–233.
- [71] J. OSKARBSKI, K. BIRR, M. MISZEWSKI, AND K. ZARSKI, Estimating the average speed of public transport vehicles based on traffic control system data, in 2015 International Conference on Models and Technologies for Intelligent Transportation Systems (MT-ITS), IEEE, 2015, pp. 287–293.
- [72] I. H. OSMAN AND J. P. KELLY, *Meta-heuristics: an overview*, in Meta-heuristics, Springer, 1996, pp. 1–21.
- [73] L. PEPPARD, String stability of relative-motion pid vehicle control systems, IEEE Transactions on Automatic Control, 19 (1974), pp. 579– 581.
- [74] M. POURABDOLLAH, E. BJÄRKVIK, F. FÜRER, B. LINDENBERG, AND K. BURGDORF, Fuel economy assessment of semi-autonomous vehicles using measured data, in 2017 IEEE Transportation Electrification Conference and Expo (ITEC), IEEE, 2017, pp. 761–766.
- [75] J. RIOS-TORRES, P. SAURAS-PEREZ, R. ALFARO, J. TAIBER, AND P. PISU, *Eco-driving system for energy efficient driving of an electric bus*, SAE International Journal of Passenger Cars-Electronic and Electrical Systems, 8 (2015), pp. 79–89.
- [76] J. RIOS-TORRES, P. SAURAS-PEREZ, R. ALFARO, J. TAIBER, AND P. PISU, Eco-driving system for energy efficient driving of an electric

bus, SAE Int. J. Passeng. Cars Electron. Electr. Syst., 8 (2015), pp. 79–89.

- [77] S. ROHKÄMPER, M. HELLWIG, AND W. RITSCHEL, Energy optimization for electric vehicles using dynamic programming, in 2017 International Conference on Research and Education in Mechatronics (REM), IEEE, 2017, pp. 1–5.
- [78] J. RYU AND J. C. GERDES, Integrating inertial sensors with global positioning system (gps) for vehicle dynamics control, J. Dyn. Sys., Meas., Control, 126 (2004), pp. 243–254.
- [79] P. SAHLHOLM AND K. H. JOHANSSON, Road grade estimation for look-ahead vehicle control using multiple measurement runs, Control Engineering Practice, 18 (2010), pp. 1328–1341.
- [80] K. SANTHANAKRISHNAN AND R. RAJAMANI, On spacing policies for highway vehicle automation, IEEE Transactions on Intelligent Transportation Systems, 4 (2003), pp. 198–204.
- [81] A. SCHWARZKOPF AND R. LEIPNIK, Control of highway vehicles for minimum fuel consumption over varying terrain, Transportation Research, 11 (1977), pp. 279–286.
- [82] Y. SEBSADJI, S. GLASER, S. MAMMAR, AND J. DAKHLALLAH, Road slope and vehicle dynamics estimation, in 2008 American control conference, IEEE, 2008, pp. 4603–4608.
- [83] D. SHEN, D. A. KARBOWSKI, AND A. ROUSSEAU, Highway ecodriving of an electric vehicle based on minimum principle, in 2018 IEEE Vehicle Power and Propulsion Conference (VPPC), Aug 2018, pp. 1–8.
- [84] Y. SUN, L. LI, B. YAN, C. YANG, AND G. TANG, A hybrid algorithm combining EKF and RLS in synchronous estimation of road grade and vehicle mass for a hybrid electric bus, Mechanical Systems and Signal Processing, 68 (2016), pp. 416–430.
- [85] D. SWAROOP AND J. HEDRICK, String stability of interconnected systems, IEEE transactions on automatic control, 41 (1996), pp. 349–357.

- [86] V. TURRI, Fuel-efficient and safe heavy-duty vehicle platooning through look-ahead control, licentiate thesis, KTH Royal Institute of Technology, 2015.
- [87] V. TURRI, B. BESSELINK, AND K. H. JOHANSSON, Cooperative lookahead control for fuel-efficient and safe heavy-duty vehicle platooning, IEEE Transactions on Control Systems Technology, PP (2016), pp. 1– 17.
- [88] C. VAGG, C. J. BRACE, D. HARI, S. AKEHURST, J. POXON, AND L. ASH, Development and field trial of a driver assistance system to encourage eco-driving in light commercial vehicle fleets, IEEE Transactions on Intelligent Transportation Systems, 14 (2013), pp. 796–805.
- [89] A. VAHIDI, A. STEFANOPOULOU, AND H. PENG, Recursive least squares with forgetting for online estimation of vehicle mass and road grade: theory and experiments, Vehicle System Dynamics, 43 (2005), pp. 31–55.
- [90] S. VAN DE HOEF, Coordination of Heavy-Duty Vehicle Platooning, PhD thesis, KTH Royal Institute of Technology, 2018.
- [91] Z. WADUD, D. MACKENZIE, AND P. LEIBY, Help or hindrance? the travel, energy and carbon impacts of highly automated vehicles, Transportation Research Part A: Policy and Practice, 86 (2016), pp. 1–18.
- [92] B. WANG, H. WANG, L. WU, L. CAI, D. PI, AND E. WANG, Truck mass estimation method based on the on-board sensor, Proceedings of the Institution of Mechanical Engineers, Part D: Journal of Automobile Engineering, 0 (0), pp. 1–15.
- [93] S. WANG, G. H. DE ALMEIDA CORREIA, AND H. X. LIN, Effects of coordinated formation of vehicle platooning in a fleet of shared automated vehicles: An agent-based model, Transportation Research Procedia, 47 (2020), pp. 377–384.
- [94] V. WINSTEAD AND I. V. KOLMANOVSKY, Estimation of road grade and vehicle mass via model predictive control, in Proceedings of 2005 IEEE Conference on Control Applications, 2005. CCA 2005., Aug 2005, pp. 1588–1593.

- [95] X. XIONG, E. XIAO, AND L. JIN, Analysis of a stochastic model for coordinated platooning of heavy-duty vehicles, arXiv preprint arXiv:1903.06741, (2019).
- [96] D. YANAKIEV AND I. KANELLAKOPOULOS, Variable time headway for string stability of automated heavy-duty vehicles, in Decision and Control, 1995., Proceedings of the 34th IEEE Conference on, vol. 4, IEEE, 1995, pp. 4077–4081.
- [97] —, Nonlinear spacing policies for automated heavy-duty vehicles, IEEE Transactions on Vehicular Technology, 47 (1998), pp. 1365–1377.
- [98] J. ZHANG AND P. A. IOANNOU, Longitudinal control of heavy trucks in mixed traffic: environmental and fuel economy considerations, IEEE Transactions on Intelligent Transportation Systems, 7 (2006), pp. 92– 104.
- [99] W. ZHANG, M. SUNDBERG, AND A. KARLSTRÖM, Platoon coordination with time windows: an operational perspective, Transportation Research Procedia, 27 (2017), pp. 357–364.
- [100] J. ZHAO, J. ZHANG, AND B. ZHU, Development and verification of the tire/road friction estimation algorithm for antilock braking system, Mathematical Problems in Engineering, 2014 (2014).
- [101] H. ZHENG, S. MA, AND Y. LIU, Vehicle braking force distribution with electronic pneumatic braking and hierarchical structure for commercial vehicle, Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering, 232 (2018), pp. 481–493.