

POLITECNICO
DI TORINO

this: POLITECNICO DI TORINO

Master Degree course in Automotive Engineering - Autonomous Driving

Master Degree Thesis

Multi-objective Adaptive Cruise Control System Design and Simulation

Supervisors
Prof. Carlo NOVARA

Candidate
Kareem FAKHREDDINE

Academic Year 2025-2026

Abstract

Modern Adaptive Cruise Control (ACC) systems are evolving from comfort features into multi-objective Advanced Driver Assistance Systems (ADAS) that must balance active safety, car-following accuracy and energy efficiency. This thesis quantifies what such systems actually gain over a realistic human driver, by simulating a two-vehicle platoon over a stochastic Torino urban driving cycle on a 4-Degree-of-Freedom (4-DOF) longitudinal plant that includes engine, clutch, gearbox, tyre slip and braking dynamics. The follower is driven either by a human driver model – an Improved Intelligent Driver Model (IIDM) with an explicit one-second reaction delay – or by one of two ADAS controllers: a Variable Time Gap (VTG) feedback law or a linear Model Predictive Controller (MPC). Both ADAS controllers command a desired acceleration that is realised by a common inner feedforward-plus-PI acceleration-tracking loop, so that the comparison isolates the outer control law. Each configuration is run over five paired driving-cycle seeds and over two powertrains, an Internal Combustion Engine (ICE) evaluated through a torque-scaled GT-POWER fuel map and an Electric Vehicle (EV) energy model with limited regenerative braking. In Case 1 (continuous operation) the results show that the benefit of the ADAS is dominated by safety rather than by energy: the minimum time-to-collision rises from 1.16 ± 0.05 s under human control to 2.90 ± 0.01 s (VTG) and 2.20 ± 0.08 s (MPC), while the harshest one-second braking falls from -7.27 ± 0.19 m/s^2 to -4.06 ± 0.36 m/s^2 (VTG) and -5.02 ± 0.87 m/s^2 (MPC), all at a comparable median time gap of 1.77 to 1.88 s. The associated energy savings are real but an order of magnitude smaller than reported in earlier versions of this study: on the EV powertrain VTG consumes $1.69 \pm 0.17\%$ and MPC $0.61 \pm 0.35\%$ less energy than the human driver, and on the ICE powertrain VTG saves $0.77 \pm 0.12\%$ while MPC saves nothing that can be distinguished from zero. A dedicated diagnostic shows that the ICE outcome is governed by the engine operating point rather than by the driving style, and in particular by the gear-shift schedule: under a speed-only shift schedule the same smooth ADAS trajectories cruise one gear lower and lose their advantage entirely. In Case 2 (dynamic safety intervention) the driver becomes distracted shortly before an emergency deceleration hazard, modelled as an increase of the reaction delay to 2.0 s. Without assistance the distracted driver collides with the leading vehicle in all five seeds, on average 62 s after the onset of the distraction; with a risk-triggered takeover to either controller, no collision occurs in any seed, and the platoon is stabilized at a minimum gap of 3.11 ± 0.16 m (VTG) or 3.21 ± 0.23 m (MPC). The VTG law also recovers more safely than the predictive controller, retaining 2.09 ± 0.42 s of minimum time-to-collision after the handover against 1.65 ± 0.69 s, while braking less hard. Overall, this work shows that the value of longitudinal automation on an urban cycle lies in disturbance attenuation and collision-risk reduction, and that its energy benefit is small and strongly conditioned by powertrain and transmission modelling assumptions.

List of Abbreviations

4-DOF	4 Degrees of Freedom
ACC	Adaptive Cruise Control
ADAS	Advanced Driver Assistance Systems
AS-RTI	Advanced-Step Real-Time Iteration
AV	Autonomous Vehicle
BMEP	Brake Mean Effective Pressure
BSFC	Brake Specific Fuel Consumption
CACC	Cooperative Adaptive Cruise Control
CAVs	Connected and Automated Vehicles
CKF	Cubature Kalman Filter
CTG	Constant-Time Gap
DHSA	Discrete Hybrid Stochastic Automata
EACC	Energy-optimal Adaptive Cruise Control
Eco-CACC	Eco-Cooperative Cruise Control
ECR	Equal Concern for Relaxation
EV	Electric Vehicle
FACMPC	Fuzzy Adaptive Cruise Control with Model Predictive Control
FLS	Fuzzy Logic System
GA	Genetic Algorithms
ICE	Internal Combustion Engine
IDM	Intelligent Driver Model
IEVs	Intelligent Electric Vehicles
IIDM	Improved Intelligent Driver Model
ITS	Intelligent Transportation Systems
KAN	Kolmogorov-Arnold Network
LMI	Linear Matrix Inequality
MD	Measured Disturbance
MIQP	Mixed-Integer Quadratic Programming
MLD	Mixed Logical Dynamical
MPC	Model Predictive Controller
MV	Manipulated Variable
NARX	Nonlinear Autoregressive with External Input
NL-CPCC	Nonlinear Predictive Cruise Control
NLP	Nonlinear Programming

NMPC	Nonlinear Model Predictive Control
PCC	Predictive Cruise Control
PI	Proportional-Integral
PID	Proportional-Derivative-Integral
PKE	Positive Kinetic Energy
RDP	Rolling Dynamic Programming
RMS	Root Mean Square
SE-MPC	Stochastic Energy Model Predictive Control
SI	Spark Ignition
SM-NMPC	Set Membership NMPC
SoC	State of Charge
SQP	Sequential Quadratic Programming
THW	Time Headway
TTC	Time To Collision
VTG	Variable Time Gap

Contents

List of Figures 6

List of Tables 8

1 Introduction 9

2 System Model 11
2.1 Overview of the Leader-Follower Architecture . . . 11
2.2 High-Fidelity Longitudinal Vehicle Dynamics . . . 12
2.2.1 Corrections applied to the inherited plant . . . 14
2.3 Transmission and Gear-Shift Strategy . . . 15
2.4 Human Driver Model . . . 15
2.5 Selected ADAS Option A: Variable Time Gap (VTG) Control . . . 17
2.6 Selected ADAS Option B: Model Predictive Control (MPC) . . . 17
2.7 Inner Acceleration-Tracking Loop . . . 19
2.8 Decision Layer . . . 20
2.9 Dual Powertrain Energy Consumption Models . . . 21
2.9.1 Internal Combustion Engine (ICE) Model . . . 21
2.9.2 Electric Vehicle (EV) Model . . . 22
2.10 Stochastic Torino Driving Cycle . . . 22
2.11 Evaluation Metrics . . . 25

3 Related Work 27
3.1 Pan et al. (2022) — Energy-Optimal MPC with Speed Prediction . . . 27
3.2 Mosharafian and Velni (2023) — Discrete Hybrid Stochastic Automata CACC . . . 28
3.3 Zhang et al. (2023) — Hierarchical Hybrid Optimal MPC for IEVs . . . 28
3.4 Boggio et al. (2023) — Fast Data-Aided Set Membership NMPC . . . 29
3.5 Mehraban et al. (2024) — Fuzzy Adaptive Cruise Control with MPC . . . 29
3.6 El-Baklish et al. (2025) — H_∞-Based Variable Time Gap Strategy . . . 30
3.7 Mosharafian and Velni (2025) — Event-Triggered Robust Tube-Based MPC 30
3.8 Li et al. (2025) — Stochastic Energy MPC with Trajectory Band Constraints 31
3.9 Dong et al. (2025) — Computationally Efficient Smooth PCC for EVs . . 31
3.10 Wang et al. (2026) — Cloud-Vehicle Collaborative Nonlinear PCC . . . 32

4 My Main Contribution 33
4.1 Experimental and Numerical Evaluation Setup . . . 33
4.2 Verification and Tuning Campaigns . . . 33
4.3 Diagnostic Methodology for the Energy Results . . . 35
4.4 Case Study Formulations and Methodology . . . 35
4.4.1 Case 1: Direct Comparative Performance (Continuous Eco-Driving and Safety) . . . 35
4.4.2 Case 2: Emergency Takeover due to Driver Distraction (Dynamic Safety Intervention) . . . 35

5 Numerical results 38
5.1 Case 1: Continuous Eco-Driving and Safety Results . . . 38
5.1.1 Safety and Comfort . . . 38
5.1.2 Energy Consumption . . . 39
5.1.3 Why the ICE and EV results differ: fuel phase decomposition . . . 41
5.1.4 Sensitivity to the gear-shift strategy . . . 42
5.1.5 VTG against MPC . . . 43
5.2 Case 2: Emergency Takeover and Safety Intervention Results . . . 44
5.2.1 Takeover Safety and Collision Avoidance . . . 44
5.2.2 Handover Transient and Controller Comparison . . . 47

6 Conclusion 49
6.1 Concluding Remarks . . . 49
6.2 Limitations . . . 50
6.3 Future Work . . . 51

Bibliography 53

List of Figures

2.1 Schematic representation of the 4-DOF longitudinal vehicle dynamics model, illustrating the geometric parameters, dynamic vertical loads, and active forces. 12

2.2 Human driver model: the perceived ego speed, leader speed and gap are delayed by the reaction time τ_r before entering the IIDM function, whose desired speed is the prevailing speed limit. 16

2.3 Implementation of the VTG control law: the variable headway, the spacing error and the saturated acceleration command. 18

2.4 MPC block and its inputs: measured outputs, references, the filtered leader acceleration as measured disturbance, the applied acceleration fed back through `ext.mv`, and the speed limit as a time-varying output bound `ymax`. 20

2.5 Decision Layer: the human and ADAS acceleration commands are routed to the inner loop according to the engagement signal, which is also logged as the ADAS status. 21

2.6 Conventional powertrain: engine speed, engine torque and vehicle speed are mapped to a fuel flow rate, which is integrated over the cycle. The function applies the torque scaling of the reference map, the idle flow and the fuel cut. 23

2.7 Electrified powertrain: the road-load model gives the mechanical power at the wheels, which is converted into battery power with constant drive and regeneration efficiencies, a regeneration power limit and a low-speed fade, and then integrated into the battery energy. 24

4.1 The simulation driving environment for one seed: leader velocity profile with the three emergency decelerations marked, and the regulatory speed limit of each zone including the anticipation ramps that precede a reduction of the limit. 34

4.2 Case 2 implementation: (a) the ego speed, leader speed and gap are delayed by either the attentive or the distracted reaction time, selected at t_{switch}, and (b) the supervisor that evaluates time-to-collision and time headway and produces the engagement flag. 37

5.1 Case 1 safety results: (a) minimum time-to-collision and harshest one-second deceleration over the five seeds, and (b) speed and spacing traces of the three controllers during one emergency deceleration hazard. 40
5.2 Paired energy differences with respect to the human follower, per seed, for both powertrains. Negative values indicate a saving of the automated controller. 41
5.3 Cruise engine speed at equal vehicle speed and resulting fuel comparison under the two gear-shift schedules. 44
5.4 One seed of Case 2: the distracted driver without assistance closes the gap to contact, while the same run with the risk-triggered takeover recovers. Magenta marks the distraction onset, red the engagement of the ADAS, and the cross the collision. The two runs are identical until the engagement, so the unassisted trace is visible only over the last seconds. 46
5.5 Case 2 outcomes over the five seeds: minimum gap after the distraction onset, where the unassisted runs reach contact, and minimum time-to-collision after the takeover. Dots are individual seeds. 47

List of Tables

2.1 Human driver model (IIDM with reaction delay) parameters. 17
2.2 VTG controller parameters. 18

5.1 Case 1 safety, headway and driving-style indicators. Mean ± standard deviation over five paired seeds; identical for the ICE and EV configurations. 38
5.2 Case 1 energy results. Mean ± standard deviation over five paired seeds. The last column is the paired difference with respect to the human follower of the same seed. 40
5.3 Fuel decomposition by driving phase (seed 5, ICE, after the 30 s warm-up). E_w is the positive work delivered at the wheels. 42
5.4 Effect of the gear-shift schedule on the ICE energy comparison (paired differences with respect to the human follower, five seeds). 43
5.5 Case 2 results after the distraction onset. Mean ± standard deviation over five paired seeds. For the unassisted runs the distracted phase extends to the collision, so the indicators before the takeover and those after the onset describe the same window. 45

Chapter 1

Introduction

The rapid proliferation of Autonomous Vehicles (AVs) and Advanced Driver Assistance Systems (ADAS) is fundamentally reshaping the automotive industry and intelligent transportation systems (ITS). Among these emerging technologies, Adaptive Cruise Control (ACC) has established itself as one of the most mature and widely commercialized ADAS features on modern roads. While traditional ACC systems were developed primarily as comfort features designed to track a constant cruising speed or maintain a fixed time headway with a preceding vehicle, modern transportation demands have forced a paradigm shift. Today's intelligent vehicles must balance a highly complex, multi-objective trade-off: guaranteeing strict active longitudinal safety, ensuring passenger comfort, and optimizing vehicle energy efficiency.

Achieving this multi-objective balance is highly challenging in dynamic and unpredictable urban environments. Conventional ACC systems utilizing rigid, fixed-spacing control logic often struggle to adapt to sudden traffic disturbances, leading to string instability (the upstream propagation and amplification of speed perturbations), compromised safety margins, or highly inefficient, rapid acceleration and braking cycles. Furthermore, accurately quantifying energy – whether the fuel consumption of an Internal Combustion Engine (ICE) or the battery energy and regenerative braking of an Electric Vehicle (EV) – requires a vehicle model that reproduces the nonlinear behaviour of the powertrain, and not only the kinematics of the trajectory.

A second, equally important requirement concerns the baseline against which automation is measured. If the human driver is represented by a poorly tuned linear regulator, any automated controller will appear dramatically superior, and the resulting figures describe the deficiencies of the baseline rather than the merits of the automation. For this reason the manual baseline used in this thesis is a calibrated car-following model taken from the traffic-flow literature – the Improved Intelligent Driver Model (IIDM) of Treiber and Kesting [15] – augmented with an explicit reaction delay [2]. This model reproduces the two behaviours that dominate manual car-following on an urban cycle: a finite perception-reaction time, and a bounded, comfort-limited acceleration response.

The primary objective of this study is to conduct an end-to-end comparative evaluation of human car-following performance against an automated ADAS framework under two distinct testing paradigms.

1. **Case 1: Direct Comparative Performance (Continuous Eco-Driving and Safety):** the follower vehicle is controlled for the whole cycle either by the human driver model or by the selected ADAS controller (VTG or MPC). This quantifies the continuous safety margins, the smoothness and the energy consumption of manual versus automated longitudinal control under identical conditions.

2. **Case 2: Emergency Takeover due to Driver Distraction (Dynamic Safety Intervention):** the simulation begins under human control, and at a randomized instant t_{switch} in the second half of the cycle (1800 s $< t_{switch} <$ 3600 s) the driver becomes distracted, modelled as a substantial increase of the reaction delay. A supervisory Decision Layer detects the resulting loss of safety margin and transfers control to the ADAS controller.

To make the comparison statistically meaningful, the driving cycle is generated by a seeded stochastic generator: the sequence of speed zones, the random target speeds inside each zone, the red traffic lights and the emergency deceleration hazards differ from seed to seed, while all controllers within one seed face exactly the same leader trajectory. Every comparison reported in this thesis is therefore a paired comparison over five seeds and over two powertrains (ICE and EV).

The contributions of this thesis are:

- a two-vehicle longitudinal simulation framework in which a reaction-delayed human driver model, a VTG law and a linear MPC drive the same 4-DOF plant through the same inner acceleration-tracking loop, so that the outer control law is the only difference between configurations;

- a paired, multi-seed evaluation protocol with safety, comfort and energy indicators (time-to-collision, time headway, peak and percentile deceleration, acceleration RMS, positive kinetic energy, and energy per 100 km) reported with dispersion and significance;

- a diagnostic decomposition of ICE fuel consumption by driving phase, which identifies the engine operating point – and specifically the gear-shift schedule – as the mechanism that governs whether smooth automated driving saves fuel at all;

- a corrected and documented plant: several defects inherited from the original coursework model (a mis-scaled gear-shift signal, a speed clamp, an always-active ABS state machine and an over-coarse solver step) are identified, fixed and reported, so that the numerical results can be reproduced.

The remainder of this document is structured as follows. Chapter 2 presents the system model: the plant, the human driver model, the two ADAS controllers, the inner loop and the two powertrain energy models, together with the driving-cycle generator and the evaluation metrics. Chapter 3 positions the study within the state of the art. Chapter 4 details the experimental methodology, the tuning campaigns and the verification tests. Chapter 5 presents the quantitative results. Chapter 6 summarizes the conclusions, states the limitations and outlines future work.

Chapter 2

System Model

To evaluate human driving against automated cruise controllers, a simulation environment must reproduce not only the trajectory of the vehicles but also the physical mechanisms through which a trajectory becomes fuel or battery energy. This chapter presents the complete physical and control architecture of the two-vehicle system.

The system is designed around a leader-follower vehicle platoon, where both vehicles are subjected to realistic external resistances, powertrain limits, and tyre-road interactions. By avoiding simplified, linear point-mass assumptions, the model captures the transient behaviour – tyre slip, clutch engagement, dynamic weight transfer, gear shifting and rotational inertia variation – that determines the engine operating point and therefore the energy result. The entire system is implemented in MATLAB/Simulink and is initialized from a parameter script and driven by batch scripts that sweep controllers, seeds and powertrains.

The longitudinal plant of both vehicles is adapted from the 4-DOF vehicle model used in the Automotive Engineering coursework at Politecnico di Torino, whose parameter set follows Genta [13]. The contribution of this thesis concerns the control architecture, the driver model, the energy models, the cycle generator and the evaluation protocol built around that plant, together with the corrections to the plant documented in Section 2.2.1.

The leader vehicle operates as the dynamic reference target: it tracks the speed profile of the selected driving cycle through a velocity controller and the same inner acceleration-tracking loop used by the follower. The follower vehicle incorporates a dual-control pathway where the active command is selected by the Decision Layer. Under Case 1 a static selector routes either the human driver model or one ADAS controller to the inner loop for the entire 3600-second run. Under Case 2 the routing changes dynamically after the driver distraction event at t_{switch}.

The remainder of the chapter gives a mathematical and structural exposition of the simulation framework developed in MATLAB/Simulink.

2.1 Overview of the Leader-Follower Architecture

The simulation environment is organized into a two-vehicle platoon consisting of a Lead Vehicle (Vehicle 1) and a Follower Vehicle (Vehicle 2, the Ego vehicle).

- **The Lead Vehicle:** the leader tracks the velocity profile of the driving cycle by means of a proportional velocity controller whose output is a desired acceleration, saturated to the same comfort and friction limits as the follower. The leader has no knowledge of the follower and no ADAS function: its role is to generate a realistic, non-smooth reference trajectory including traffic-light stops and severe deceleration hazards.

- **The Follower Vehicle:** the follower is the subject of the comparative study. It contains a human driver subsystem and an ADAS subsystem (VTG or MPC), both producing a desired acceleration, and a Decision Layer that selects which of the two reaches the inner acceleration-tracking loop [Figure 2.5].

Because the human driver model and the two ADAS controllers all command the same physical quantity (a desired longitudinal acceleration) and share the same inner loop and the same plant, every difference observed in the results is attributable to the outer control law alone.

The Simulink implementation of each of the three candidate controllers is shown in the section that describes it: the human driver model in Figure 2.2, the VTG law in Figure 2.3 and the MPC in Figure 2.4; the Decision Layer that routes one of them to the inner loop is shown in Figure 2.5.

2.2 High-Fidelity Longitudinal Vehicle Dynamics

The plant is a 4-Degree-of-Freedom (4-DOF) longitudinal dynamic model [Figure 2.1] whose states are the vehicle speed, the engine speed and the two wheel speeds. The vehicle is a compact passenger car of mass $m = 1020$ kg, with a naturally aspirated spark-ignition engine of maximum torque $T_{e,max} = 105$ N $\cdot$ m, a five-speed gearbox with ratios $i_g = [3.909,\ 2.157,\ 1.480,\ 1.121,\ 0.897]$, a final drive ratio $i_f = 3.438$ and an effective rolling radius $R_e = 0.279$ m.

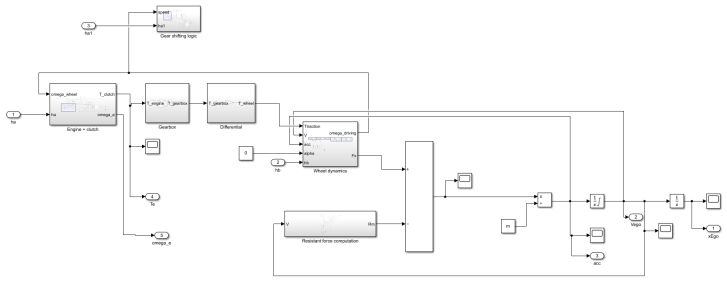


Figure 2.1: Schematic representation of the 4-DOF longitudinal vehicle dynamics model, illustrating the geometric parameters, dynamic vertical loads, and active forces.

- **Vehicle Translation (Velocity, V):** the translational motion of the vehicle mass m is governed by the balance of traction forces at the driving wheels, friction forces

at the non-driving wheels, and environmental resistance:

$$m\dot{V} = \eta_f F_{x1} + \eta_r F_{x2} - R_m(V) - mg\sin\alpha \tag{2.1}$$

where V is the longitudinal velocity, η_f and η_r are the number of front and rear wheels, F_{x1} and F_{x2} are the tractive forces at the front and rear contact patches, α is the road grade (zero in all simulations of this thesis) and $R_m(V)$ is the running resistance polynomial:

$$R_m(V) = A + BV^2 + CV^4 \tag{2.2}$$

with

$$\begin{aligned} A &= mgf_0, \\ B &= mgK + \tfrac{1}{2}\rho S(C_x - C_z f_0), \\ C &= -\tfrac{1}{2}\rho S K C_z, \end{aligned} \tag{2.3}$$

where f_0 and K are the rolling resistance coefficients, S is the frontal area, ρ the air density and C_x, C_z the drag and lift coefficients. The same polynomial is reused as the feedforward road-load model of the inner control loop (Section 2.7), which guarantees that the controller and the plant share an identical description of the resistances.

- **Engine/Motor Rotation (ω_e):** the rotational speed of the power source is governed by the balance between the delivered torque T_e and the clutch load torque T_c:

$$J_e\dot{\omega_e} = T_e - T_c \tag{2.4}$$

where J_e is the combined inertia of engine and flywheel. The engine torque follows a normalized throttle command $h_a \in [0,1]$ scaled by the full-load torque curve:

$$T_e = h_a\, T_{e,max}(\omega_e) \tag{2.5}$$

- **Front Driving Wheels (ω_{w1}):**

$$J^*_{w1}\dot{\omega_{w1}} = \eta_t \frac{T_c}{\tau_g\tau_f} - T_{b1} - F_{x1}R_e \tag{2.6}$$

where T_{b1} is the front braking torque, $\tau_g = 1/i_g$ and $\tau_f = 1/i_f$, and J^*_{w1} is the equivalent inertia of the front axle, which depends on the engaged gear:

$$J^*_{w1} = \frac{J_c}{(\tau_g\tau_f)^2} + \frac{J_t}{\tau_f^2} + \sum J_{wi} \tag{2.7}$$

The clutch torque T_c follows a piecewise law depending on whether the clutch is slipping, locked, or transmitting engine braking:

$$T_c = \begin{cases} h_c T_{cmax}, & \text{if } \Omega_e > \frac{\Omega_{w1}}{\tau_g\tau_f} \quad \text{(Slipping)} \\ h_a T_{emax}, & \text{if } \Omega_e = \frac{\Omega_{w1}}{\tau_g\tau_f} \quad \text{(Synchronized)} \\ -h_c T_{cmax}, & \text{if } \Omega_e < \frac{\Omega_{w1}}{\tau_g\tau_f} \quad \text{(Engine Braking)} \end{cases} \tag{2.8}$$

- **Rear Non-Driving Wheels (ω_{w2}):** the rear wheels receive no drive torque and are driven by the vehicle motion, counteracted by the rear braking torque T_{b2}:

$$J_{w2}\dot{\omega_{w2}} = -F_{x2}R_e - T_{b2} \tag{2.9}$$

- **Tyre forces:** the tractive forces are computed from a longitudinal friction curve evaluated at the slip ratio σ_i of each axle,

$$F_{xi} = \mu(\sigma_i)\, F_{zi} \tag{2.10}$$

where $\mu(\sigma)$ is a lookup table for a dry asphalt surface, with its peak near $\sigma \approx 0.15$. The vertical loads capture the longitudinal weight transfer:

$$F_{z1} = \frac{mgl_r - ma_x h_{cg}}{l_f + l_r}, \quad F_{z2} = \frac{mgl_f + ma_x h_{cg}}{l_f + l_r} \tag{2.11}$$

- **Braking:** the friction brakes are modelled by a front torque saturation $T_{bf,max} = 800\ \mathrm{N \cdot m}$ and a rear torque obtained through the brake distribution ratio K_b. The brake command $h_b \in [0,1]$ scales these limits.

2.2.1 Corrections applied to the inherited plant

Four defects of the inherited model were identified during verification and corrected; they are documented here because they directly affect the reported quantities.

1. **Gear-shift signal scaling.** The speed signal entering the gear-shift logic was scaled by $1/(R_e\tau_f)$ instead of by the final-drive ratio i_f, which mis-scaled the shift points by a factor of about 3.6. The gain was corrected to i_f and the shift thresholds were set to physically meaningful engine speeds.

2. **Speed clamp.** The velocity integrator of both vehicles carried a lower saturation of 0.8 m/s, which prevented the vehicles from stopping and made every standstill metric meaningless. The limits were set to $[0, 25]$ m/s and the initial condition to the cycle initial speed.

3. **Brake saturation and ABS.** The brake actuator saturation was raised to the physical limit $T_{bf,max}$, and the inherited ABS state machine – which clamped the brake torque to a fraction of a newton-metre during normal driving – was bypassed, since wheel lock does not occur within the deceleration range used in this study.

4. **Solver step.** With the original fixed step of 10^{-2} s the wheel-speed derivative became non-finite near standstill once the speed clamp was removed. The fixed-step size was reduced to $T_s = 10^{-3}$ s with an `ode4` solver, which removes the artefact.

2.3 Transmission and Gear-Shift Strategy

The gearbox is operated by an automated shift logic that selects the gear from the engine speed $\omega_e = \omega_t\, i_g(k)$, where ω_t is the gearbox output speed. A purely speed-based schedule with fixed thresholds was found to interact strongly with the driving style (Section 5.1.4), and was replaced by a throttle-dependent schedule that reproduces the qualitative behaviour of a production transmission controller: early upshifts at light throttle, late upshifts at full throttle, and kickdown on demand. The upshift threshold is interpolated linearly with the throttle command h_a,

$$n_{up}(h_a) = n_{up,eco} + (n_{up,full} - n_{up,eco})\, h_a \tag{2.12}$$

with $n_{up,eco} = 2000$ rpm, $n_{up,full} = 2800$ rpm, and the downshift threshold fixed at $n_{down} = 1300$ rpm. An upshift from gear k to $k+1$ is accepted only if the resulting engine speed remains safely above the downshift threshold,

$$\omega_t\, i_g(k+1) > 1.1\, n_{down} \tag{2.13}$$

which prevents shift hunting across the wide first-to-second ratio step. A kickdown to gear $k-1$ is commanded when $h_a > 0.9$ and the resulting engine speed stays below $0.9\, n_{up,full}$.

2.4 Human Driver Model

The manual baseline is the Improved Intelligent Driver Model (IIDM) of Treiber and Kesting [15], a car-following model widely used in traffic-flow research and calibrated against empirical trajectory data. Compared with the original IDM, the improved formulation removes the unrealistic widening of the equilibrium gap when the driver approaches its desired speed, which is essential on an urban cycle where the vehicle spends much of its time close to the speed limit.

Let s be the bumper-to-bumper gap, v the ego speed, v_L the leader speed, $\Delta v = v - v_L$ the approaching rate and v_0 the desired speed (the prevailing speed limit). The desired dynamical gap is

$$s^*(v, \Delta v) = s_0 + \max\!\left(0,\; vT + \frac{v\,\Delta v}{2\sqrt{ab}}\right) \tag{2.14}$$

and, with $z = s^*/s$, the acceleration is

$$a = \begin{cases} a\left(1 - z^2\right), & v \le v_0,\; z \ge 1 \\ a_f\left(1 - z^{2a/a_f}\right), & v \le v_0,\; z < 1 \\ a_f + a\left(1 - z^2\right), & v > v_0,\; z \ge 1 \\ a_f, & v > v_0,\; z < 1 \end{cases} \tag{2.15}$$

where the free-flow term is

$$a_f = \begin{cases} a\left[1 - \left(\dfrac{v}{v_0}\right)^{\delta}\right], & v \le v_0 \\ -b\left[1 - \left(\dfrac{v_0}{v}\right)^{a\delta/b}\right], & v > v_0 \end{cases} \tag{2.16}$$

Human perception and actuation are not instantaneous. Following Treiber et al. [2], all three perceived quantities – ego speed, leader speed and gap – are delayed by a constant reaction time τ_r before entering the model, which reproduces the delayed, and therefore amplifying, response of a manual driver to a braking leader. The resulting acceleration is finally saturated to the same comfort and friction bounds used by the ADAS controllers. The Simulink implementation is shown in Figure 2.2: the three perceived signals pass through Transport Delay blocks before entering the model, while the desired speed is the regulatory limit of the current zone.

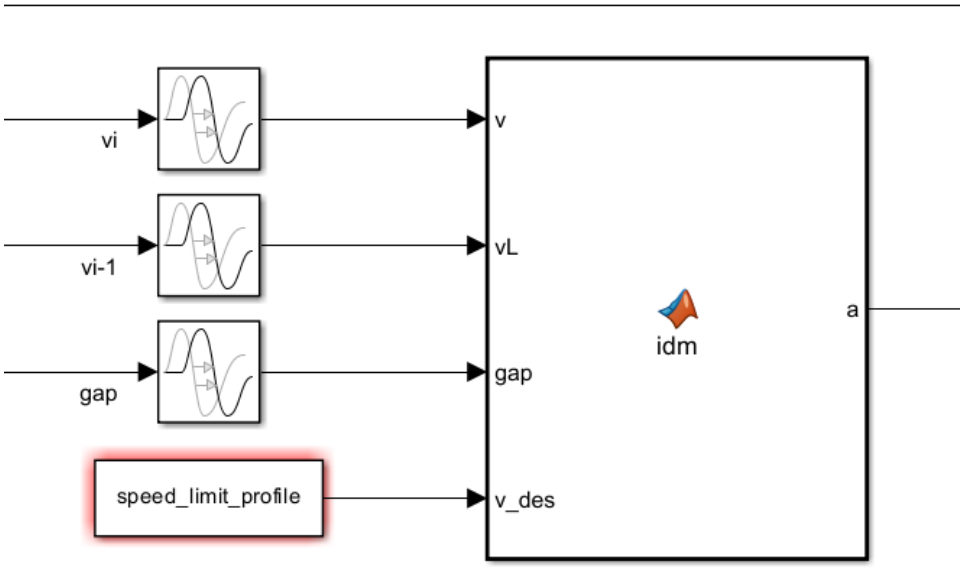


Figure 2.2: Human driver model: the perceived ego speed, leader speed and gap are delayed by the reaction time τ_r before entering the IIDM function, whose desired speed is the prevailing speed limit.

The parameters are listed in Table 2.1; they are the standard urban values recommended in [15], with a reaction time of 1.0 s representative of an attentive driver responding to an expected event.

Table 2.1: Human driver model (IIDM with reaction delay) parameters.

Symbol	Description	Value
T	desired time gap	1.5 s
s_0	minimum standstill gap	3.0 m
a	comfortable acceleration	1.0 m/s^2
b	comfortable deceleration	1.5 m/s^2
δ	free acceleration exponent	4
τ_r	reaction delay	1.0 s
a_{min}	friction-limited deceleration	$-0.8g$
a_{max}	acceleration saturation	1.0 m/s^2

2.5 Selected ADAS Option A: Variable Time Gap (VTG) Control

The first automated controller is a spacing-error feedback law with a velocity-dependent headway. The baseline is the constant-time-gap policy of Rajamani [14], in which the desired spacing grows linearly with the ego speed, extended by the variable-headway modulation of Yanakiev and Kanellakopoulos [1], which shortens the headway when the gap is opening and lengthens it when the gap is closing. This modulation is what gives the law its disturbance-attenuating behaviour during braking transients, in the same spirit as the $\mathcal{H}_\infty$-based strategy of El-Baklish et al. [9].

With $s = x_{lead} - x_{ego} - L$ the bumper-to-bumper gap, L the vehicle length and $\Delta v = v_{lead} - v_{ego}$ the relative speed, the headway is

$$h = \mathrm{sat}(h_0 - c_h \Delta v,\ h_{min},\ h_{max}) \tag{2.17}$$

the spacing error with respect to the resulting policy is

$$\delta = s - (s_0 + h\, v_{ego}) \tag{2.18}$$

and the commanded acceleration is

$$a_{des} = \mathrm{sat}\left(\frac{\Delta v + \lambda\, \delta}{h_0},\ a_{min},\ a_{max}\right) \tag{2.19}$$

The gain λ sets the convergence rate of the spacing error; h_0 normalizes the loop so that the closed-loop response is independent of the instantaneous headway. The control law is implemented as the algebraic function of Figure 2.3, evaluated at every solver step.

The parameters are given in Table 2.2; the value $\lambda = 0.8$ s^{-1} was selected from the sweep described in Section 4.2.

2.6 Selected ADAS Option B: Model Predictive Control (MPC)

The second automated controller is a linear Model Predictive Controller that optimizes the future acceleration profile over a finite horizon subject to explicit constraints. The

```
function [acc, h, gap] = vtg_law(x_ego, v_ego, x_prec, v_prec, vtg)
gap   = (x_prec - x_ego) - vtg.L;                    % bumper-to-bumper gap
dv    = v_prec - v_ego;                              % > 0: leader pulling away
h     = min(max(vtg.h0 - vtg.c_h*dv, vtg.h_min), vtg.h_max);
delta = gap - (vtg.s0 + h*v_ego);                    % spacing error
acc   = (dv + vtg.lambda*delta) / vtg.h0;
acc   = min(max(acc, vtg.a_min), vtg.a_max);
end
```

Figure 2.3: Implementation of the VTG control law: the variable headway, the spacing error and the saturated acceleration command.

Table 2.2: VTG controller parameters.

Symbol	**Description**	**Value**
h_0	nominal time gap	1.5 s
c_h	headway modulation gain	0.1 s^2/m
h_{min}, h_{max}	headway bounds	1.0, 2.5 s
s_0	standstill spacing	3.0 m
λ	spacing error gain	0.8 s^{-1}
L	vehicle length	4.5 m
a_{min}, a_{max}	acceleration bounds	$-0.8g$, 1.0 $\mathrm{m/s}^2$

prediction model is a continuous-time state-space representation of the headway dynamics augmented with a first-order actuator lag, so that the controller is aware of the finite bandwidth of the inner acceleration-tracking loop:

$$\dot{x}(t) = Ax(t) + B_u u(t) + B_v v(t) \tag{2.20}$$

with the state vector $x(t) = [\, v_{ego}(t);\ s(t);\ v_{lead}(t);\ a_{ego}(t) \,]^T$, the manipulated variable $u(t) = a_{des}(t)$ and the measured disturbance $v(t) = a_{lead}(t)$:

$$A = \begin{bmatrix} 0 & 0 & 0 & 1 \\ -1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1/\tau_a \end{bmatrix},\ B_u = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 1/\tau_a \end{bmatrix},\ B_v = \begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} \tag{2.21}$$

where $\tau_a = 0.1$ s is the identified time constant of the inner loop (Section 4.2). The controlled outputs are the ego speed and the spacing margin relative to a constant time gap $h_{mpc} = 1.5$ s:

$$C = \begin{bmatrix} 1 & 0 & 0 & 0 \\ -h_{mpc} & 1 & 0 & 0 \end{bmatrix}, \quad \begin{aligned} y_1 &= v_{ego} \\ y_2 &= s - h_{mpc} v_{ego} \end{aligned} \tag{2.22}$$

The reference of y_1 is the prevailing speed limit of the current zone, supplied as a time-varying reference signal, and the reference of y_2 is the standstill spacing $s_0 = 3$ m, so that in steady state the controller settles on the spacing policy $s = s_0 + h_{mpc} v_{ego}$ rather than on a zero-gap target.

The model is discretized with a sampling period $T_{s,mpc} = 0.1$ s; the prediction horizon is $P = 40$ steps (4.0 s) and the control horizon is $M = 6$ steps. At each instant the controller minimizes

$$\begin{aligned} J(k) = \sum_{j=1}^{P} \Big(& w_{y_1}[y_1(k+j|k) - r_1(k+j)]^2 \\ & + w_{y_2}[y_2(k+j|k) - r_2(k+j)]^2 \Big) \\ & + \sum_{j=0}^{M-1} \Big(w_u[u(k+j|k)]^2 + w_{\Delta u}[\Delta u(k+j|k)]^2 \Big) \end{aligned} \tag{2.23}$$

with the weights $w_{y_1} = 0.2$, $w_{y_2} = 1.0$, $w_u = 0.1$ and $w_{\Delta u} = 1.0$, selected by the sweep of Section 4.2. The spacing output is weighted five times more than the speed output, so that headway regulation dominates speed-limit tracking; the rate weight penalizes jerk and produces smooth acceleration profiles.

The optimization is subject to

$$-0.8g \leq u(k) \leq 1.0 \text{ m/s}^2 \tag{2.24}$$

$$-1.0 \text{ m/s}^2 \leq \Delta u(k) \leq 0.5 \text{ m/s}^2 \tag{2.25}$$

$$y_2(k) \geq 2.0 \text{ m} \tag{2.26}$$

$$y_1(k) \leq v_{limit}(k) \tag{2.27}$$

The input bounds represent the friction limit in braking and a comfort limit in traction; the rate bounds limit jerk. The spacing constraint is softened with an Equal Concern for Relaxation weight ECR = 0.1, so that the quadratic program remains feasible during severe hazards, when no admissible acceleration can preserve the nominal margin. The speed-limit constraint is enforced as a time-varying output upper bound rather than as a fixed constant, which allows the controller to anticipate a reduction of the limit at a zone boundary.

Two implementation details are relevant for the comparison. First, the measured disturbance is the leader acceleration filtered by a first-order lag of 0.1 s; a heavier filter delays the controller response to hazards by approximately half a second and measurably reduces the minimum time-to-collision. Second, the controller receives the acceleration actually applied to the plant through its optional manipulated-variable input, which provides bumpless transfer when control is handed over – a property that is essential in Case 2. Both features are visible in the block implementation of Figure 2.4, together with the measured-disturbance filter and the time-varying speed-limit bound.

2.7 Inner Acceleration-Tracking Loop

The three outer controllers produce a desired acceleration; they do not produce pedal commands. A common inner loop converts a_{des} into the throttle and brake commands of the plant. It combines a physical feedforward term with a proportional-integral correction on the acceleration error, and runs at 100 Hz.

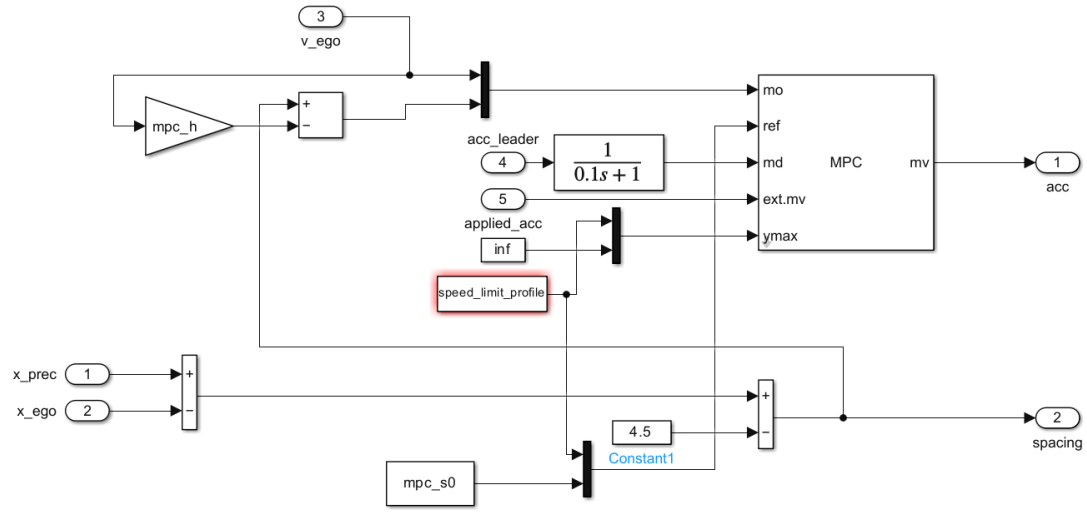


Figure 2.4: MPC block and its inputs: measured outputs, references, the filtered leader acceleration as measured disturbance, the applied acceleration fed back through `ext.mv`, and the speed limit as a time-varying output bound `ymax`.

The required longitudinal force follows from the inverse of the road-load model:

$$F_{req} = m\,a_{des} + A + Bv^2 + Cv^4 \tag{2.28}$$

and is normalized by the maximum available tractive or braking force to produce the feedforward command

$$u_{ff} = \begin{cases} F_{req}/F_{drive}, & F_{req} \geq 0 \\ F_{req}/F_{brake}, & F_{req} < 0 \end{cases} \tag{2.29}$$

The realized acceleration is estimated by differentiating the measured speed and filtering it with a first-order filter of time constant $\tau_f = 0.05$ s; using this estimate instead of the raw plant acceleration avoids aliasing of the high-frequency driveline oscillations that occur near standstill. The command is

$$u = u_{ff} + K_p\,(a_{des} - \hat{a}) + I, \qquad u_s = \mathrm{sat}(u, -1{,}1) \tag{2.30}$$

with conditional integration (the integrator is frozen when the command is saturated and the error would drive it further into saturation), $K_p = 0.1$ and $K_i = 0.5$. Finally the throttle and brake commands are $h_a = \max(u_s{,}0)$ and $h_b = \max(-u_s{,}0)$. Below 0.5 m/s with a non-positive acceleration request, the loop ramps a small holding brake pressure, which reproduces the behaviour of a driver holding the vehicle at a traffic light and prevents creeping.

2.8 Decision Layer

The Decision Layer routes the active acceleration command to the inner loop.

- **Case 1 (Continuous Comparative Mode):** the layer acts as a static switch configured before the run. With the ADAS disabled the vehicle is driven by the human model for the entire cycle; with the ADAS enabled the human model is bypassed and the selected controller (VTG or MPC) drives for the entire cycle.

- **Case 2 (Dynamic Safety Intervention Mode):** the simulation starts under human control. At t_{switch} the driver becomes distracted, which is modelled by an increased reaction delay. A supervisor monitors the time-to-collision and the time headway computed from the undelayed (sensor) signals and engages the ADAS when a risk threshold is crossed or when a driver-monitoring detection delay has elapsed, whichever occurs first, plus an actuation latency. Because the switch acts on the desired acceleration upstream of the inner loop, and because the MPC is initialized with the acceleration currently applied, the handover is bumpless.

The Simulink implementation of the Decision Layer is shown in Figure 2.5. In Case 1 the routing signal is a constant set before the run; in Case 2 it is the engagement flag produced by the supervisor of Figure 4.2b.

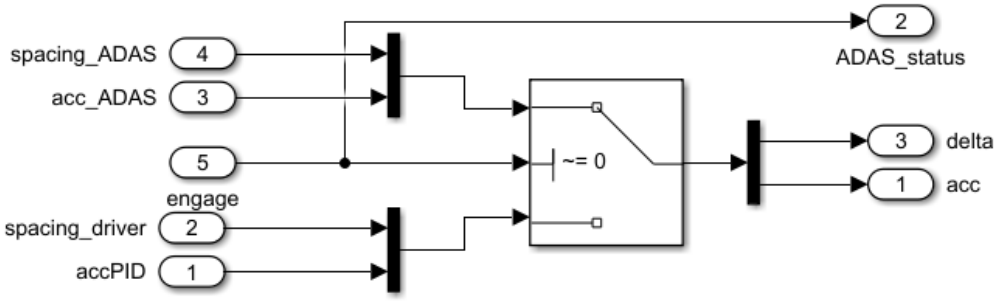


Figure 2.5: Decision Layer: the human and ADAS acceleration commands are routed to the inner loop according to the engagement signal, which is also logged as the ADAS status.

2.9 Dual Powertrain Energy Consumption Models

2.9.1 Internal Combustion Engine (ICE) Model

The conventional powertrain converts the instantaneous engine speed ω_e and engine torque T_e produced by the 4-DOF plant into a fuel mass flow through a steady-state consumption map [Figure 2.6]:

$$m_f = \int_0^{T_{sim}} \dot{m}_f(\omega_e, T_e)\, dt \tag{2.31}$$

The map is built from the publicly available GT-POWER engine dataset distributed with the MATLAB Powertrain Blockset documentation (a naturally aspirated 2.0 L spark-ignition inline-four), from which the best brake-specific fuel consumption point of each speed sweep is extracted and interpolated over a regular grid in engine speed and torque. Because the simulated engine is smaller than the mapped one, the operating point is transferred at equal brake mean effective pressure through the torque ratio

$$k_T = \frac{T_{max,map}}{T_{max,sim}} = \frac{179.5}{104} = 1.726 \tag{2.32}$$

so that the fuel flow of the simulated engine is $\dot{m}_f(\omega_e, T_e) = \dot{m}_{f,map}(\omega_e, k_T T_e)/k_T$. Two operating modes complete the model: an idle consumption of 0.124 g/s (approximately 0.6 L/h) whenever the vehicle is stationary, and a fuel cut whenever the torque demand vanishes while the vehicle is moving above 1 m/s. The resulting map exhibits brake-specific values of approximately 430 to 450 g/kWh at light load and 255 to 300 g/kWh at medium and high load, which is the load dependence that later proves decisive for the interpretation of the ICE results.

2.9.2 Electric Vehicle (EV) Model

The electrified powertrain is an energy-accounting model rather than a component-level motor and battery model [Figure 2.7]. The mechanical power at the wheels is converted into battery power through constant drive and regeneration efficiencies, with an auxiliary load $P_{aux} = 400$ W:

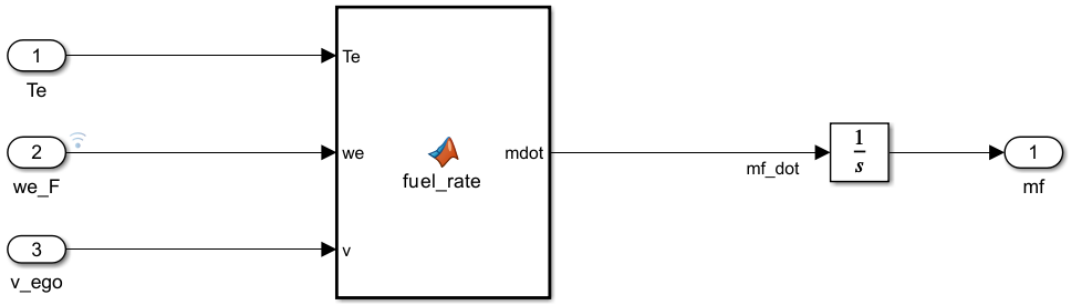
$$P_{batt} = \begin{cases} \dfrac{P_{mech}}{\eta_{drive}} + P_{aux}, & P_{mech} \geq 0 \\ \eta_{regen}\, \varphi(v) \max\left(P_{mech}, -P_{regen,max}\right) + P_{aux}, & P_{mech} < 0 \end{cases} \tag{2.33}$$

with $\eta_{drive} = 0.90$, $\eta_{regen} = 0.85$ and a regeneration power limit $P_{regen,max} = 50$ kW. The factor $\varphi(v) = \min(v/v_{fade}, 1)$ with $v_{fade} = 5$ km/h fades regeneration out at very low speed, where an electric machine can no longer recover energy and the friction brakes take over. Integrating P_{batt} gives the energy drawn from a 60 kWh battery initialized at a state of charge of 0.8.

It must be stressed that the EV and the ICE configurations share the same chassis, the same mass and the same transmission: the EV results therefore quantify the effect of the driving style on an electrified energy path, and are not a model of a specific production electric vehicle. The absence of the additional battery mass, of temperature effects and of speed-dependent motor efficiency explains why the absolute consumption figures are lower than those of commercial electric vehicles of comparable size.

2.10 Stochastic Torino Driving Cycle

The leader trajectory is produced by a seeded generator that emulates an urban itinerary through Torino. The route is divided into thirteen zones of duration between 60 and 720 s. Each zone is characterized by a speed band whose bounds lie between 20 and

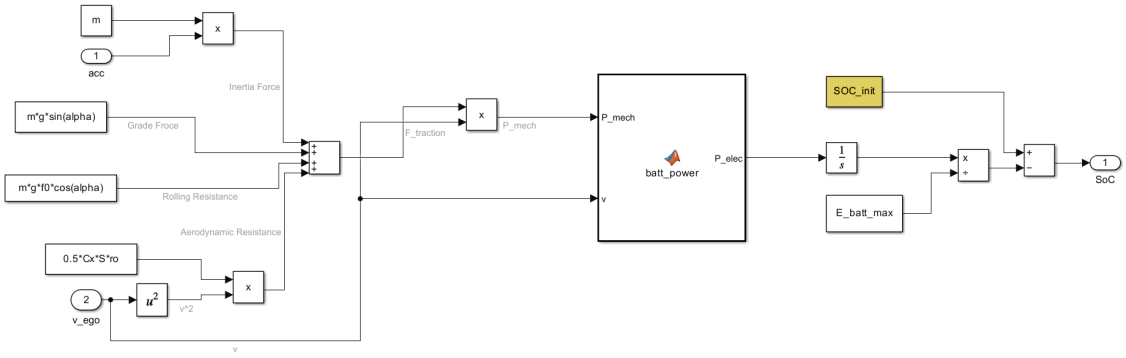


(a) Subsystem

```
function mdot = fuel_rate(Te, we, v, fm)
% Fuel mass flow [g/s]: GT-POWER map scaled to this engine, with idle and fuel cut.
if Te <= fm.Te_off
    if v > fm.v_cut
        mdot = 0;                    % overrun (coasting in gear): fuel cut
    else
        mdot = fm.idle;              % standing or creeping: idle
    end
    return
end
rpm  = min(max(we*30/pi, fm.rpm(1)), fm.rpm(end));
Teq  = min(Te*fm.kT, fm.Tq(end));
mdot = interp2(fm.rpm, fm.Tq, fm.tbl, rpm, Teq) * 1000 / fm.kT;   % kg/s (2.0 L) -> g/s (vehicle)
mdot = max(mdot, fm.idle);           % a running engine never uses less than idle
end
```

(b) Fuel rate function

Figure 2.6: Conventional powertrain: engine speed, engine torque and vehicle speed are mapped to a fuel flow rate, which is integrated over the cycle. The function applies the torque scaling of the reference map, the idle flow and the fuel cut.



(a) Subsystem

```
function P_elec = batt_power(P_mech, v, ev)
% Battery power [W] from mechanical power at the wheels.
if P_mech >= 0
    P_elec = P_mech/ev.eta_drive + ev.P_aux;              % traction
else
    fade   = min(v/ev.v_fade, 1);                         % no regen near standstill
    P_rgn  = max(P_mech, -ev.P_regen_max) * fade;         % motor regen power limit
    P_elec = P_rgn*ev.eta_regen + ev.P_aux;               % remainder -> friction brakes
end
end
```

(b) Battery power function

Figure 2.7: Electrified powertrain: the road-load model gives the mechanical power at the wheels, which is converted into battery power with constant drive and regeneration efficiencies, a regeneration power limit and a low-speed fade, and then integrated into the battery energy.

70 km/h, and the upper bound of the band is taken as the regulatory speed limit of the zone, so that the limits encountered along the route range from 30 to 70 km/h. Inside each zone a target speed is drawn every 10 s within the band and interpolated smoothly; each of the traffic lights of a zone is red with probability 0.2 and, when red, imposes a 20 s standstill. The resulting profile is rate-limited to $\pm 2.0/2.5$ m/s^2 and smoothed.

Three emergency hazards are injected at random instants, each consisting of a -6.0 m/s^2 deceleration to standstill, a 5 s stop and a controlled recovery ramp. These events are the dominant safety stimulus of the cycle and are the situations in which the difference between a delayed human response and an automated response becomes visible.

The generator also returns the regulatory speed limit as a function of time, which is used both as the desired speed of the human driver model and as the speed reference and upper bound of the MPC. A backward pass enforces that a reduction of the limit is anticipated with a comfortable deceleration of 1.5 m/s^2,

$$v_{lim}(k) \leftarrow \min\left(v_{lim}(k),\ v_{lim}(k+1) + a_{lim}\,\Delta t\right) \tag{2.34}$$

Without this anticipation, a step reduction of the limit at a zone boundary makes the MPC output constraint instantaneously infeasible and provokes an artificial emergency braking manoeuvre that has no counterpart in real driving.

Each seed produces a different but statistically equivalent cycle; within one seed, all controllers and both powertrains face exactly the same leader trajectory, which enables paired statistical comparisons. The driving cycle and the corresponding speed zones are illustrated in Figure 4.1.

2.11 Evaluation Metrics

All metrics are computed after a warm-up of 30 s, which removes the initial transient in which the platoon settles from its initial conditions into its steady spacing policy.

Safety. The minimum time-to-collision, evaluated only while the follower is closing on the leader,

$$\mathrm{TTC}(t) = \frac{s(t)}{v_{ego}(t) - v_{lead}(t)}, \qquad v_{ego} > v_{lead} \tag{2.35}$$

is the primary safety indicator, complemented by the minimum bumper-to-bumper gap, the median standstill gap, and the fraction of moving time spent with a time headway below 1 s.

Comfort and severity. The harshest sustained braking is measured as the minimum of the one-second moving average of the follower acceleration, which avoids attributing severity to numerical spikes; the 95^{th} percentile deceleration and the fraction of time below -3.5 m/s^2 describe the rest of the distribution.

Driving style. The acceleration RMS and the positive kinetic energy per unit distance,

$$\mathrm{PKE} = \frac{\sum_k \left(v_{k+1}^2 - v_k^2\right)^+}{D} \tag{2.36}$$

quantify how aggressively the vehicle accelerates. PKE is computed for both vehicles, so that the follower can be compared with the leader it follows: a follower whose PKE exceeds the leader's amplifies the disturbances it receives, while a follower below it attenuates them.

Energy. Consumption is normalized by distance and reported as litres per 100 km (ICE, using a fuel density of 745 kg/m^3) or kilowatt-hours per 100 km (EV). Two comparisons are reported: the follower against its own leader within the same run, and the follower against the human follower of the same seed.

Statistics. Every comparison is paired by seed. With five seeds, the paired differences are summarized by their mean and standard deviation and tested with a paired t-test with four degrees of freedom. A Wilcoxon signed-rank test cannot return a two-sided p-value below 0.0625 with five pairs, which is why the parametric test is used; the paired differences are reported individually so that the reader can verify their consistency.

Chapter 3

Related Work

After establishing the system model, it is essential to contextualize this architecture within the current state of the art. The following works propose advanced ACC solutions, balancing tracking performance, safety, and energy efficiency.

3.1 Pan et al. (2022) — Energy-Optimal MPC with Speed Prediction

Pan et al. [3] propose an energy-optimal adaptive cruise control (EACC) strategy for pure electric vehicles based on a Model Predictive Control (MPC) framework.

- *Scenario & Contribution:* They integrate a simplified motor energy consumption model directly as the economic evaluation index in the MPC quadratic cost function. To improve speed tracking and energy recovery during high-speed deceleration, a Nonlinear Autoregressive with External Input (NARX) neural network is developed to predict the acceleration profile of the leading vehicle, which is treated as an external disturbance in the MPC loop. Additionally, a variable weight strategy is introduced to dynamically trade off safety-related spacing tracking at low speeds and economy-maximizing torque smoothing at high speeds.

- *Comparison to this Thesis:* Both works address the multi-objective trade-off between longitudinal tracking accuracy and powertrain efficiency. However, Pan et al. implement a continuous variable-weight strategy based on real-time speed thresholds. In contrast, this thesis keeps the controller structure fixed and instead varies the driver: the same plant is driven by a calibrated human car-following model and by two automated controllers, which isolates the benefit attributable to automation itself. Furthermore, this thesis integrates a 4-DOF chassis plant containing clutch and gear-shifting logic, and evaluates both an ICE and an EV energy path under an identical, non-smooth urban speed profile.

3.2 Mosharafian and Velni (2023) — Discrete Hybrid Stochastic Automata CACC

Mosharafian and Velni [5] investigate the problem of Cooperative Adaptive Cruise Control (CACC) for connected and automated vehicles (CAVs) under intermittent communication constraints and distance measurement errors.

- *Scenario & Contribution:* They model a CACC system using a discrete hybrid stochastic automata (DHSA) framework. The control system dynamically switches among three discrete operating modes: free following, warning, and emergency braking, formulated in a Mixed Logical Dynamical (MLD) form. By solving a mixed-integer quadratic programming (MIQP) problem online, the controller minimizes spacing deviations while explicitly penalizing unnecessary emergency braking transitions to avoid destabilizing the vehicle platoon.

- *Comparison to this Thesis:* Both studies recognize the necessity of an active safety trigger to handle emergency deceleration hazards. However, Mosharafian and Velni design a hybrid controller that continuously switches operating states online based on probability weights and V2V communication states. This thesis, conversely, uses a static routing in Case 1 to measure the performance bounds of manual versus automated control – an increase of the minimum time-to-collision from 1.16 s to 2.90 s with a simultaneous reduction of peak braking from -7.27 to -4.06 m/s^2 – and reserves dynamic switching for the driver-distraction scenario of Case 2.

3.3 Zhang et al. (2023) — Hierarchical Hybrid Optimal MPC for IEVs

Zhang et al. [6] propose a hierarchical control framework designed to synchronize active safety, mobile energy storage management, and fuel economy in Intelligent Electric Vehicles (IEVs).

- *Scenario & Contribution:* The upper-level master controller employs a constrained hybrid optimal MPC algorithm, combining Genetic Algorithms (GA) for global parameter search with Sequential Quadratic Programming (SQP) to handle non-linear constraints and optimize energy flow. The lower-level inner loop uses hysteresis current control to regulate the drive/generating motor and battery storage system, ensuring precise longitudinal force tracking via an electronic longitudinal control system.

- *Comparison to this Thesis:* While both works model electric powertrain dynamics, battery energy and regenerative braking during deceleration, Zhang et al. focus on complex hierarchical solvers (GA-SQP) and hysteresis current control for active force distribution. In contrast, this thesis designs a computationally lighter, self-contained ADAS (either MPC or VTG) running on standard solvers and expands the comparative evaluation to include an internal combustion powertrain evaluated

through a torque-scaled GT-POWER fuel map under a stochastic stop-and-go Torino driving cycle.

3.4 Boggio et al. (2023) — Fast Data-Aided Set Membership NMPC

Boggio et al. [4] tackle the real-time computational limitations of Nonlinear Model Predictive Control (NMPC) in autonomous driving applications.

- *Scenario & Contribution:* They develop a Set Membership NMPC (SM-NMPC) approach for trajectory planning and control. The system uses offline data-driven Set Membership approximation to derive tight, guaranteed bounds on the optimal control law. These bounds are applied online to restrict the search domain of the SQP solver, providing a robust "warm start" initialization that reduces the online NMPC optimization time by approximately six times in sinusoidal lane-keeping and parallel parking maneuvers.
- *Comparison to this Thesis:* Both systems seek to overcome the numerical latency of non-linear control to ensure real-time implementation on physical hardware. However, while Boggio et al. focus on restricting the search space of multi-variable NMPCs (including steering and lateral dynamics) using offline data learning, this thesis implements a linear MPC on a short prediction horizon and shows that, on an urban car-following task, its measurable advantage over a well-tuned algebraic feedback law is limited – a result that is itself an argument about where predictive complexity is worth its cost.

3.5 Mehraban et al. (2024) — Fuzzy Adaptive Cruise Control with MPC

Mehraban et al. [7] investigate the problem of adapting longitudinal cruise controllers to rapidly changing, non-linear traffic conditions.

- *Scenario & Contribution:* They present a Fuzzy Adaptive Cruise Control with Model Predictive Control (FACMPC) system. The algorithm utilizes a Fuzzy Logic System (FLS) to dynamically modulate the weighting parameters $(\alpha_1, \alpha_2, \beta)$ of the MPC quadratic cost function in real-time based on the instantaneous spacing error, relative velocity, and ego acceleration. This adaptive tuning enhances the controller's robustness against environmental disturbances (such as variations in road slope, wind speed, and rolling resistance) without the extensive training requirements of neural networks.
- *Comparison to this Thesis:* Both studies optimize the conflicting objectives of comfort, safety, and energy efficiency. While Mehraban et al. rely on a fuzzy inference engine to continuously modulate the controller priorities, this thesis fixes the weights through an explicit sensitivity sweep and shows that, over a full urban cycle,

the energy outcome is far more sensitive to powertrain and transmission modelling assumptions than to the controller weights themselves.

3.6 El-Baklish et al. (2025) — H_∞-Based Variable Time Gap Strategy

El-Baklish et al. [9] address the trade-off between traffic capacity and string stability in ACC systems.

- *Scenario & Contribution:* They propose a Variable Time Gap (VTG) control strategy utilizing a nonlinear $\mathcal{H}_\infty$ state feedback policy. Traditional constant time-gap (CTG) settings can amplify traffic disturbances upstream through a platoon, causing string instability and localized traffic jams. The proposed VTG controller dynamically adjusts the target time headway based on transient vehicle states, allowing the vehicle to maintain tight spacing during steady-state cruising while proactively increasing the safety gap to absorb sudden braking perturbations.

- *Comparison to this Thesis:* This research directly inspires the VTG control law implemented in our selectable ADAS framework. However, while El-Baklish et al. evaluate string stability and traffic throughput using microscopic traffic simulators, this thesis embeds the VTG controller into a 4-DOF longitudinal chassis model with two powertrain energy paths. This makes it possible to confirm the attenuation property physically – the VTG follower reproduces the positive kinetic energy of its leader (0.224 against 0.222 $\mathrm{m/s^2}$) whereas the human follower amplifies it (0.235 $\mathrm{m/s^2}$) – and to convert that attenuation into fuel and battery energy.

3.7 Mosharafian and Velni (2025) — Event-Triggered Robust Tube-Based MPC

Mosharafian and Velni [11] address the challenges of limited wireless communication bandwidth and computational resources in Cooperative ACC (CACC) platoons.

- *Scenario & Contribution:* They propose an event-triggered, tube-based robust MPC framework. Instead of transmitting data and solving the MPC problem periodically at fixed intervals, the control task is initiated only when specific tracking error or safety thresholds are violated. A linear matrix inequality (LMI) condition is solved offline to guarantee platoon stability, allowing the system to achieve a 70% reduction in communication frequency compared to conventional periodic systems while still ensuring the vehicles reach their desired speeds and following distances.

- *Comparison to this Thesis:* Both control systems seek to optimize execution efficiency while ensuring robust safety. However, Mosharafian and Velni focus on a cooperative, multi-agent network that reduces communication bandwidth through a continuous nominal-and-error feedback tube architecture. This thesis, conversely, implements a localized, standalone ADAS architecture that evaluates a single follower

under strictly human or strictly automated control, with zero V2V communication dependency.

3.8 Li et al. (2025) — Stochastic Energy MPC with Trajectory Band Constraints

Li et al. [10] develop a stochastic energy model predictive control (SE-MPC) framework designed to optimize the energy efficiency of Eco-Cooperative Adaptive Cruise Control (Eco-CACC) systems under stochastic traffic conditions.

- *Scenario & Contribution:* They integrate a Kolmogorov-Arnold Network (KAN) for highly accurate preceding vehicle speed prediction and a Cubature Kalman Filter (CKF) for robust state estimation. To prevent computational failures and solver infeasibility under uncertainty, they introduce a "trajectory band loss function". This approach replaces rigid "hard" constraints with probabilistic soft boundaries scaled by the state covariance matrix, improving motor operating efficiency by 0.8% and reducing NEDC energy consumption by 5.14% compared to classical MPC.
- *Comparison to this Thesis:* Both works integrate predictive algorithms with MPC to enhance energy efficiency in unpredictable traffic. The energy improvements reported here are of a similar order of magnitude to the motor-efficiency gain reported by Li et al., which supports the conclusion that single-vehicle longitudinal smoothing yields percent-level rather than tens-of-percent energy savings. This thesis additionally quantifies how much of such a gain survives the transition from an electrified to a conventional powertrain.

3.9 Dong et al. (2025) — Computationally Efficient Smooth PCC for EVs

Dong et al. [8] address the real-time execution challenges of Predictive Cruise Control (PCC) in autonomous electric vehicles.

- *Scenario & Contribution:* They identify a critical numerical hurdle in PCC: the nonsmooth terms in vehicle dynamics, such as the sudden efficiency changes and mathematical discontinuities when the powertrain switches between motor traction and regenerative braking. They propose an equivalent smooth mathematical reformulation that transforms the non-smooth optimal control problem into a smooth Nonlinear Programming (NLP) problem without introducing modeling simplifications. The resulting NLP is solved via an Advanced-Step Real-Time Iteration (AS-RTI) scheme, reducing HIL energy consumption by 1% within a 50 ms computation limit.
- *Comparison to this Thesis:* Both studies capture longitudinal EV dynamics including regeneration. While Dong et al. eliminate nonsmooth model transitions mathematically to accelerate NLP solvers, this thesis retains the physical discontinuities – gear

shifts, clutch states, fuel cut and regeneration fade – inside the plant, and demonstrates that these discontinuities, rather than the control algorithm, dominate the final energy figure.

3.10 Wang et al. (2026) — Cloud-Vehicle Collaborative Nonlinear PCC

Wang et al. [12] present a cloud-based Nonlinear Predictive Cruise Control (NL-CPCC) framework to optimize fuel efficiency for heavy-duty commercial vehicles on highways.

- *Scenario & Contribution:* To overcome the computational limitations of modeling highly non-linear vehicle dynamics (including slope-inertia coupling and engine torque efficiency maps) in real-time, they propose a cloud-vehicle collaborative architecture. A cloud-based rolling dynamic programming (RDP) algorithm generates long-term, globally optimal reference trajectories and feasible warm-start initializations. The vehicle edge then executes a short-horizon, real-time nonlinear predictive control algorithm to ensure safety and comfort, achieving fuel savings of up to 2.40%.
- *Comparison to this Thesis:* Both works recognize that linear simplifications compromise true energy optimization, and both embed nonlinear longitudinal dynamics into their evaluation. However, Wang et al. distribute the computational burden between cloud and vehicle edge for continuous highway cruising, where long-horizon terrain preview is the dominant source of savings. This thesis addresses the opposite regime – dense urban stop-and-go traffic with a short preview horizon – and finds savings of comparable magnitude but arising from a different mechanism, namely the attenuation of the leader disturbances.

Chapter 4

My Main Contribution

This chapter details the control strategies, the configuration logic, the tuning campaigns and the comparative evaluation methodology developed for this thesis. The contribution of this work is a simulation framework in which a realistic human driver and two automated controllers can be compared on strictly equal terms, together with the diagnostic procedures needed to explain – and not merely to report – the resulting energy figures.

4.1 Experimental and Numerical Evaluation Setup

The platoon is simulated in Simulink with a fixed-step `ode4` solver at $T_s = 10^{-3}$ s, in accelerator mode; the logged signals are sampled at 100 Hz. Each run covers 3600 s of the stochastic Torino cycle described in Section 2.10.

The evaluation campaign of Case 1 consists of three controllers (human, VTG, MPC) $\times$ five seeds $\times$ two powertrains (ICE, EV), that is thirty runs. Within a seed, the leader trajectory, the traffic lights, the hazards and the speed-limit profile are identical for all controllers, and the leader consumption is verified to be identical across the three runs, which validates the pairing. The profile of the Torino driving cycle and its speed zones are illustrated in Figure 4.1.

4.2 Verification and Tuning Campaigns

Before the comparative campaign, each element of the architecture was verified in isolation.

Inner loop identification. The follower plant was excited with a sequence of acceleration steps ($+0.5$, -2, 0 and -6 m/s^2). The closed inner loop tracks the requested acceleration with a 63% response time of approximately 0.11 s in traction and 0.01 to 0.04 s in braking, and with steady-state errors below 0.05 m/s^2 except at the friction limit, where the requested -6 m/s^2 yields -6.13 m/s^2. The identified traction time constant is the value $\tau_a = 0.1$ s used in the MPC prediction model.

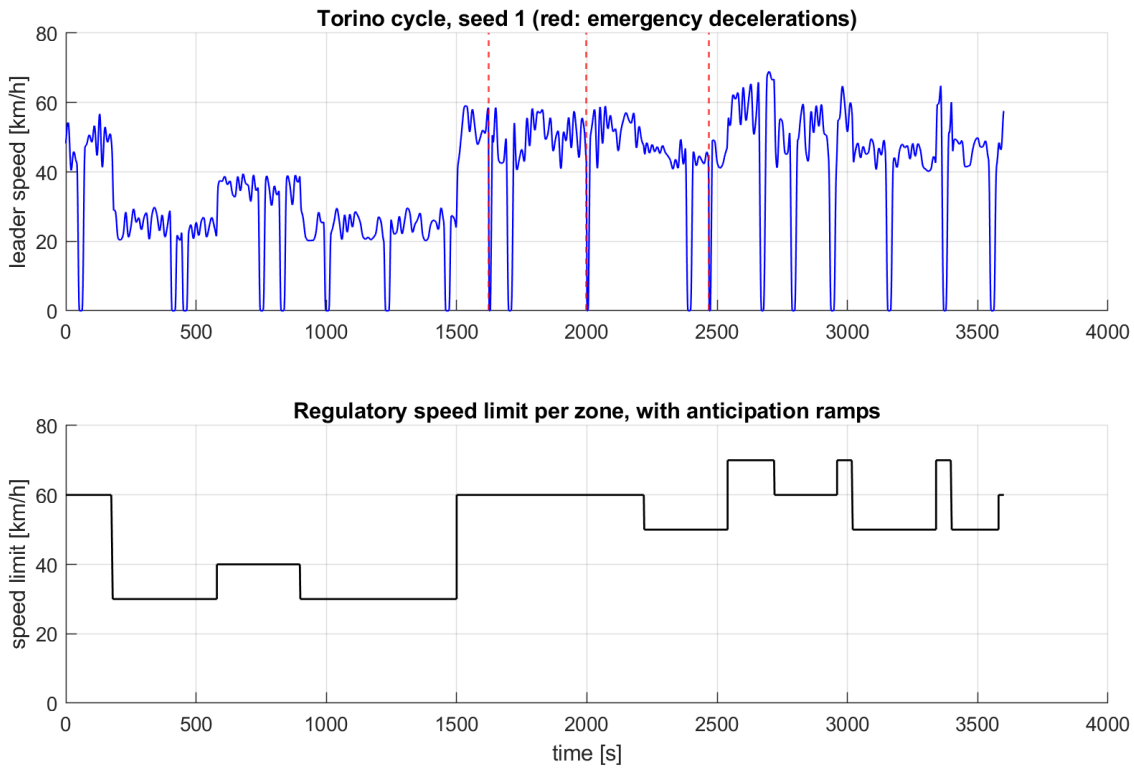


Figure 4.1: The simulation driving environment for one seed: leader velocity profile with the three emergency decelerations marked, and the regulatory speed limit of each zone including the anticipation ramps that precede a reduction of the limit.

VTG gain. The spacing-error gain λ was swept from 0.2 to 1.5 s^{-1}. Below 0.6 s^{-1} the spacing error converges too slowly and the minimum gap degrades; above 0.8 s^{-1} all indicators are essentially flat. The value $\lambda = 0.8$ s^{-1} was therefore selected as the smallest gain achieving full performance, which also minimizes control activity.

MPC weights. The output weights, the input weight and the rate weight were swept over a grid. Across the whole grid the fuel consumption varied by less than 0.6%, whereas the seed-to-seed variation of the same quantity reaches 10% to 13%: the energy result is therefore governed by the cycle, not by the weight selection. The final weights were chosen for their transient behaviour rather than for their consumption, favouring headway regulation and jerk limitation.

Tyre slip regularization. A constant slip term used in the wheel model was reduced from 1 to 0.1 to test its influence. The identified acceleration time constants were unchanged, while the number of sign changes of the tyre force near standstill dropped from 253 to 6, confirming that the term acts as a numerical regularization and not as a physical parameter.

Cycle and constraint consistency. An early version of the cycle applied the regulatory speed limit as a step at each zone boundary. With the MPC speed-limit constraint active, this produced a spurious -6.6 m/s^2 emergency stop at the first boundary of every run. The backward anticipation pass of Section 2.10 was introduced to remove this artefact, which is a modelling error rather than a controller failure.

4.3 Diagnostic Methodology for the Energy Results

Because the measured energy differences between controllers are of the order of one percent, a dedicated diagnostic was developed to establish their physical origin rather than to report them as an unexplained outcome. For a given seed, the follower fuel flow is integrated separately over five mutually exclusive driving phases – standstill, fuel cut, acceleration, cruise and fuelled deceleration – and compared with the positive mechanical work delivered at the wheels in the same phase,

$$E_w = \int \max\left(F_x v, 0\right) dt, \qquad F_x = ma + A + Bv^2 + Cv^4 \tag{4.1}$$

The ratio of fuel to wheel work in each phase gives a brake-specific consumption that can be compared between controllers, and the engine speed is logged so that the transmission state can be inspected at equal vehicle speed. This decomposition is what allows the results of Section 5.1.4 to be attributed to the gear-shift schedule rather than to the control law.

4.4 Case Study Formulations and Methodology

4.4.1 Case 1: Direct Comparative Performance (Continuous Eco-Driving and Safety)

The follower is driven for the entire cycle by one of the three controllers:

- **Human Driver Active:** the reaction-delayed IIDM of Section 2.4 controls the vehicle from $t = 0$ to $t = 3600$ s.
- **ADAS Controller Active (VTG or MPC):** the human model is bypassed and the selected controller manages headway and speed for the entire run.

The objective of Case 1 is to quantify, under identical disturbances, the safety margins, the smoothness and the energy consumption of manual against automated longitudinal control. Since all three controllers settle on a comparable median time gap (1.77 to 1.88 s) and cover the same distance to within a few metres over 37 km, the comparison is not confounded by differences in headway policy or throughput.

4.4.2 Case 2: Emergency Takeover due to Driver Distraction (Dynamic Safety Intervention)

This scenario tests the transient behaviour of the system when the driver becomes distracted:

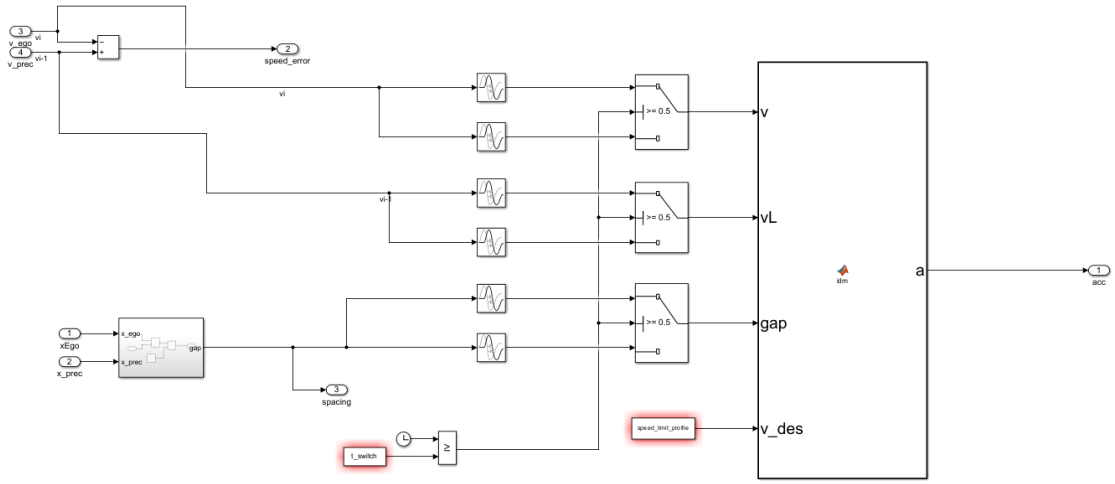
- **Initial Attentive Phase:** the follower is driven by the human model with its nominal reaction delay.
- **Temporal Switching Logic (t_{switch}):** the distraction onset is *anchored to a hazard*. The cycle generator returns the instants of the emergency decelerations it has injected, and the onset is placed a random 30 to 90 s before the latest hazard of the second half of the cycle. Without this anchoring, a uniformly distributed onset leaves the distracted phase of some seeds free of any hazard, and the scenario would measure the probability of meeting an emergency while distracted rather than the performance of the intervention itself. Each seed therefore contains exactly one distracted encounter with a -6.0 m/s^2 event, preceded by normal attentive car-following.
- **Attention Degradation:** distraction is modelled by increasing the reaction delay of the driver model from 1.0 to 2.0 s, a value representative of a distracted driver. All other parameters of the driver model are unchanged: the driver's intent is identical, only its perception is older. Two parallel delay paths are implemented and switched at t_{switch}, so that the attentive and distracted drivers are the same model.
- **Automated ADAS Takeover:** a supervisor evaluates, from the undelayed sensor signals, the time-to-collision and the time headway

$$\mathrm{TTC} = \frac{s}{v_{ego} - v_{lead}}, \qquad \mathrm{THW} = \frac{s}{v_{ego}} \tag{4.2}$$

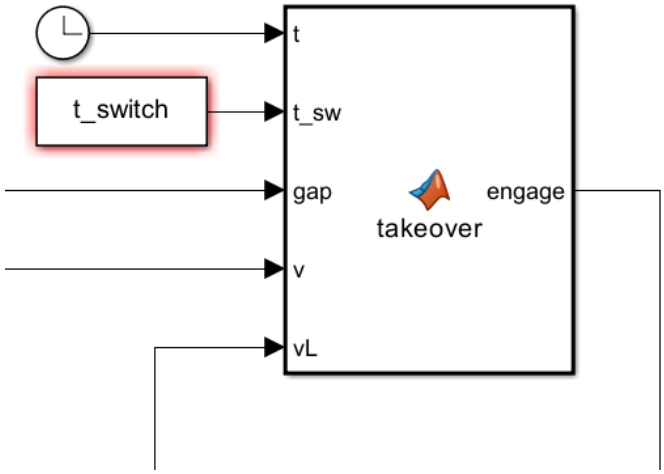
  and requests a takeover when TTC < 3.0 s or THW < 0.8 s. The ADAS is engaged after an actuation latency of 0.3 s and remains engaged for the rest of the run. The supervisor acts on the desired acceleration upstream of the inner loop, and the MPC is initialized with the acceleration currently applied, so the handover does not introduce a discontinuity beyond the manoeuvre the new controller intends.

The Simulink implementation of the distracted driver and of the supervisor is shown in Figure 4.2: the two delay chains feed the same driver model through three switches driven by the distraction flag, and the supervisor evaluates the risk from the undelayed sensor signals.

This is a purely risk-triggered policy: the supervisor does not know that the driver is distracted and reacts only to the developing danger. The reference for the scenario is the same configuration with the supervisor disabled, so that the benefit of the intervention is quantified against the distracted driver itself and not only against the attentive baseline. Three configurations are simulated over the same five seeds: distracted driver without assistance, takeover to VTG and takeover to MPC.



(a) Attentive and distracted perception paths



(b) Takeover supervisor

Figure 4.2: Case 2 implementation: (a) the ego speed, leader speed and gap are delayed by either the attentive or the distracted reaction time, selected at t_{switch}, and (b) the supervisor that evaluates time-to-collision and time headway and produces the engagement flag.

Chapter 5

Numerical results

This chapter presents the quantitative results of both simulation campaigns, Case 1 in Section 5.1 and Case 2 in Section 5.2, each over five paired seeds and both powertrains. All values are reported as mean ± standard deviation over the five seeds, and all comparisons are paired within seeds.

5.1 Case 1: Continuous Eco-Driving and Safety Results

5.1.1 Safety and Comfort

Table 5.1 collects the safety and comfort indicators. Since the plant, the cycle and the controllers are identical in the ICE and EV campaigns, the trajectories – and therefore all safety metrics – are identical in the two powertrain configurations; this was verified run by run and constitutes a consistency check on the simulation campaign.

Table 5.1: Case 1 safety, headway and driving-style indicators. Mean ± standard deviation over five paired seeds; identical for the ICE and EV configurations.

Indicator	**Human**	**VTG**	**MPC**
Minimum TTC (s)	1.16 ± 0.05	$\mathbf{2.90 \pm 0.01}$	2.20 ± 0.08
Harshest 1 s braking (m/s^2)	-7.27 ± 0.19	$\mathbf{-4.06 \pm 0.36}$	-5.02 ± 0.87
Time below -3.5 m/s^2 (%)	0.28 ± 0.13	$\mathbf{0.11 \pm 0.06}$	0.14 ± 0.06
Minimum gap (m)	2.43 ± 0.05	3.03 ± 0.01	$\mathbf{3.34 \pm 0.01}$
Standstill gap (m)	2.92 ± 0.04	3.08 ± 0.01	3.39 ± 0.01
Median time gap (s)	1.88 ± 0.03	1.79 ± 0.02	1.77 ± 0.02
Acceleration RMS (m/s^2)	0.521 ± 0.064	$\mathbf{0.456 \pm 0.050}$	0.488 ± 0.057
PKE (m/s^2)	0.235 ± 0.023	$\mathbf{0.224 \pm 0.027}$	0.234 ± 0.029
Leader PKE (m/s^2)	0.222 ± 0.023		
Collisions	0/5	0/5	0/5

The dominant result of this thesis is the safety margin. Under continuous manual

control the minimum time-to-collision falls to 1.16 ± 0.05 s, below the threshold of approximately 1.5 s commonly used to identify critical car-following situations, whereas the VTG and MPC controllers hold 2.90 ± 0.01 s and 2.20 ± 0.08 s respectively. The mechanism is visible in the braking statistics. The hazards of the cycle require the leader to brake at -6.0 m/s^2; the human follower responds with a sustained deceleration of -7.27 m/s^2, that is it *amplifies* the disturbance by approximately 20%, while VTG attenuates it to 68% and MPC to 84% of the leader deceleration. This amplification is the direct consequence of the reaction delay: the manual driver begins to brake approximately one second after the leader, and must then recover the lost distance with a harder manoeuvre.

The same ordering appears in the driving-style indicators. The human follower produces a positive kinetic energy of 0.235 m/s^2 against the 0.222 m/s^2 of the leader it follows, whereas the VTG follower reproduces the leader value (0.224 m/s^2) and falls below it in three of the five seeds. In the language of platoon dynamics, the manual driver behaves as a string-unstable element and the VTG controller as a string-stable one, and this is observed on a physical plant rather than on a kinematic model.

It is essential for the fairness of the comparison that these benefits are not obtained by following further behind: the median time gap of the human driver (1.88 s) is in fact slightly *larger* than that of the two automated controllers (1.79 and 1.77 s), and the distance covered over the cycle differs by at most 21 m in 37 km. The ADAS controllers thus achieve better safety margins at equal or tighter headway and equal throughput.

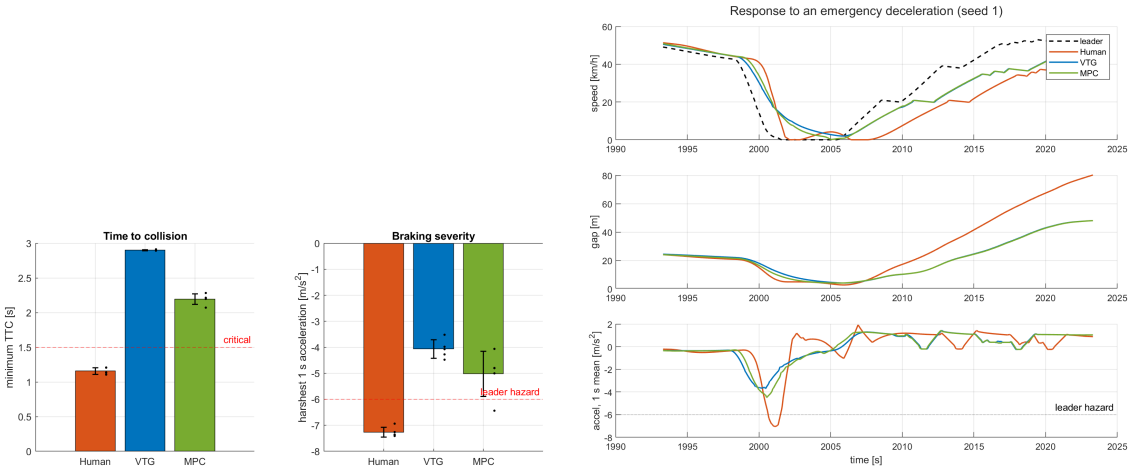
The minimum gap deserves a separate remark. For all three controllers the minimum gap occurs at standstill, not during motion, so it measures the stopping discipline of the controller rather than a near-collision: the human driver overshoots its own standstill target of 3 m and stops at 2.43 m, whereas the MPC settles at 3.34 m. The dynamic safety margin is measured by the time-to-collision, and there the difference between manual and automated control is unambiguous. Finally, the worst single event of the entire campaign was produced by the MPC in seed 4 (-6.44 m/s^2, TTC 2.07 s), which illustrates the larger dispersion of the predictive controller: its response depends on the evolution of the prediction, whereas the VTG response is an algebraic function of the current state and is consequently far more repeatable (± 0.01 s on the minimum TTC across seeds).

5.1.2 Energy Consumption

Table 5.2 reports the energy results for both powertrains.

On the EV powertrain both automated controllers consume less energy than the human driver in every one of the five seeds. The paired differences are $-1.69 \pm 0.17\%$ for VTG (paired t-test, $t = 22.2$, 4 degrees of freedom, $p < 0.001$) and $-0.61 \pm 0.35\%$ for MPC ($t = 3.9$, $p = 0.018$). The ordering of the energy results reproduces exactly the ordering of the driving-style indicators of Table 5.1: the controller that best attenuates the leader disturbances consumes the least. This is the expected physical behaviour of an electrified powertrain, whose efficiency depends only weakly on the operating point, so that every additional acceleration costs energy of which regeneration recovers only a fraction.

On the ICE powertrain the picture changes. VTG saves $0.77 \pm 0.12\%$ with respect to the human driver, consistently across the five seeds ($t = 14.0$, $p < 0.001$), but the



(a) Safety and comfort indicators

(b) Response to a deceleration hazard

Figure 5.1: Case 1 safety results: (a) minimum time-to-collision and harshest one-second deceleration over the five seeds, and (b) speed and spacing traces of the three controllers during one emergency deceleration hazard.

Table 5.2: Case 1 energy results. Mean ± standard deviation over five paired seeds. The last column is the paired difference with respect to the human follower of the same seed.

Configuration	**Follower consumption**	**vs. leader (%)**	**vs. human (%)**
ICE powertrain — L/100 km (leader: 5.575 ± 0.128)			
Human	5.671 ± 0.131	$+1.72 \pm 0.36$	—
VTG	$\mathbf{5.628 \pm 0.128}$	$+0.94 \pm 0.31$	$\mathbf{-0.77 \pm 0.12}$
MPC	5.669 ± 0.130	$+1.69 \pm 0.41$	-0.03 ± 0.26
EV powertrain — kWh/100 km (leader: 6.898 ± 0.106)			
Human	7.021 ± 0.126	$+1.77 \pm 0.37$	—
VTG	$\mathbf{6.903 \pm 0.117}$	$+0.06 \pm 0.19$	$\mathbf{-1.69 \pm 0.17}$
MPC	6.978 ± 0.127	$+1.15 \pm 0.52$	-0.61 ± 0.35

saving is less than half of its EV counterpart; the MPC saving, $-0.03 \pm 0.26\%$, cannot be distinguished from zero. Moreover, all three followers consume *more* than their own leader (+0.94% to +1.72%), although the VTG follower reproduces the leader trajectory almost exactly in kinetic terms. Car-following therefore has an intrinsic energy cost on a conventional powertrain: the corrections that keep the headway regulated are small, frequent and executed at light engine load, where a spark-ignition engine is least efficient.

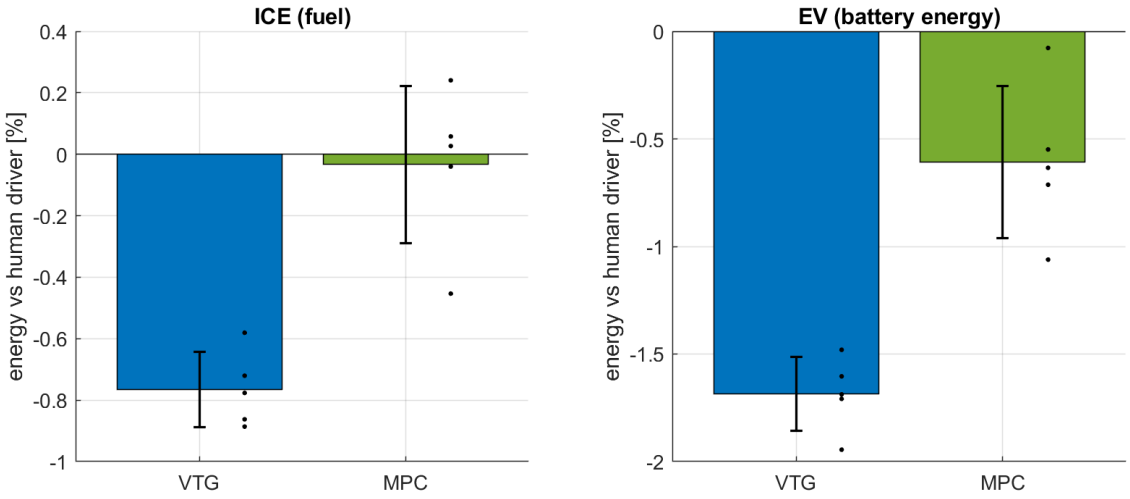


Figure 5.2: Paired energy differences with respect to the human follower, per seed, for both powertrains. Negative values indicate a saving of the automated controller.

5.1.3 Why the ICE and EV results differ: fuel phase decomposition

To identify the mechanism behind the ICE result, the diagnostic of Section 4.3 was applied to the seed with the largest discrepancy. Table 5.3 reports the decomposition for the human and VTG followers.

The decomposition shows that the VTG follower performs 4.7% less mechanical work at the wheels than the human follower (2.335 against 2.449 kWh) but converts it into only 0.5% less fuel. The reason is the redistribution of time between phases: by avoiding the accelerate-and-brake pattern of the delayed driver, the VTG follower spends 300 s more in cruise and 206 s less in acceleration. Acceleration is the efficient phase of a spark-ignition engine (403 to 419 g/kWh), while cruise is a light-load phase in which the same engine needs 797 to 837 g/kWh. Smooth driving therefore shifts operation towards the least efficient region of the engine map, and a substantial part of the mechanical saving is lost in the conversion. The EV powertrain has no such load dependence, which is why the same trajectories yield a saving twice as large there.

This also explains the apparently paradoxical observation that every follower consumes more fuel than its leader while covering the same distance at the same average speed: the follower must continuously correct its headway, and those corrections are executed in the light-load region.

Table 5.3: Fuel decomposition by driving phase (seed 5, ICE, after the 30 s warm-up). E_w is the positive work delivered at the wheels.

Phase	**Time (s)**	**Fuel (g)**	E_w **(kWh)**	**g/kWh**
Human				
Standstill	141.9	21.3	—	—
Fuel cut	65.0	0.0	—	—
Acceleration	1026.8	692.8	1.654	419
Cruise	1423.9	603.7	0.758	797
Deceleration	912.5	238.0	0.037	—
Total		**1556**	**2.449**	**635**
VTG				
Standstill	64.8	11.4	—	—
Fuel cut	60.8	0.0	—	—
Acceleration	820.4	577.2	1.433	403
Cruise	1723.3	717.2	0.857	837
Deceleration	900.7	242.2	0.045	—
Total		**1548**	**2.335**	**663**

5.1.4 Sensitivity to the gear-shift strategy

The load dependence discussed above makes the ICE result sensitive to the transmission control, and this sensitivity turned out to be strong enough to reverse the conclusion. The campaign was therefore repeated with two shift schedules: the original one, in which upshifts and downshifts depend only on engine speed (2800 and 1300 rpm), and the throttle-dependent schedule of Section 2.3.

Under the speed-only schedule the two ADAS controllers consumed *more* fuel than the human driver, and the difference varied strongly from seed to seed. The engine-speed measurement explains why. With the ratios of the simulated gearbox, third gear corresponds to 2008 rpm and fourth gear to 1521 rpm at 41.5 km/h; with a fixed upshift threshold of 2800 rpm, fourth gear can only be engaged above 57.9 km/h and is only released below 35.5 km/h. Between those two speeds the engaged gear depends on the history of the trajectory. The human driver, whose delayed response makes it overshoot the leader speed, crosses the upshift threshold and cruises in fourth gear; the smooth ADAS trajectories never reach it and cruise in third gear, at 121 rpm higher engine speed for the same vehicle speed and a lower power demand. The penalty of the higher engine speed exceeded the benefit of the smoother trajectory: at the phase-specific consumption rates of the human driver, the VTG follower would have consumed 1% *less* fuel, so the

Table 5.4: Effect of the gear-shift schedule on the ICE energy comparison (paired differences with respect to the human follower, five seeds).

Quantity	**Speed-only schedule**	**Throttle-dependent schedule**
Leader consumption (L/100 km)	6.05 to 6.46	5.43 to 5.75
VTG vs. human (%)	$+1.38 \pm 1.77$	-0.77 ± 0.12
MPC vs. human (%)	$+1.62 \pm 1.24$	-0.03 ± 0.26
Cruise engine speed, human (rpm)	1797	1641
Cruise engine speed, VTG (rpm)	1918	1646
Median cruise speed (km/h)	41.5	41.4

entire +3% observed was an artefact of the gear selection.

Replacing the schedule with a throttle-dependent one, which upshifts early at light throttle as a production transmission controller does, equalizes the cruise engine speed (1641 against 1646 rpm at the same cruise speed), reduces the absolute consumption of both vehicles by 6% to 9%, brings the leader consumption into a realistic range for a 1020 kg petrol car in urban traffic, and restores the ordering observed on the EV powertrain.

This sensitivity is reported here as a result in its own right. It shows that on a conventional powertrain the eco-driving benefit of an ACC controller cannot be established independently of the transmission control strategy, and that studies that omit the shift schedule may report the sign of the effect incorrectly.

5.1.5 VTG against MPC

Across the campaign, the VTG law outperforms the linear MPC on every indicator except the minimum standstill gap: it holds a 32% larger minimum time-to-collision, brakes 19% less harshly in the worst event, produces a 7% lower acceleration RMS, saves nearly three times more energy on the EV powertrain, and is the only controller whose ICE saving is statistically distinguishable from zero. It is also far more repeatable across seeds.

Two reasons explain this outcome. First, the MPC minimizes a quadratic cost over a four-second horizon using a linear model of the headway dynamics; on a cycle dominated by discrete events – traffic lights, speed-limit changes and unpredicted hazards – that horizon is too short to anticipate the events that matter and long enough for the linear prediction to be inaccurate during the transients that follow. Second, the leader acceleration reaches the controller through a measured-disturbance filter, so the predictive controller effectively responds to hazards slightly later than the algebraic feedback law, which acts on the instantaneous state. The practical implication is that, for this task, the additional complexity of the predictive controller is not repaid: a well-tuned variable-headway feedback law achieves better safety and better energy performance while requiring only the evaluation of an algebraic expression at each step, against the solution of a constrained quadratic program at 10 Hz. The solver time was not measured in this

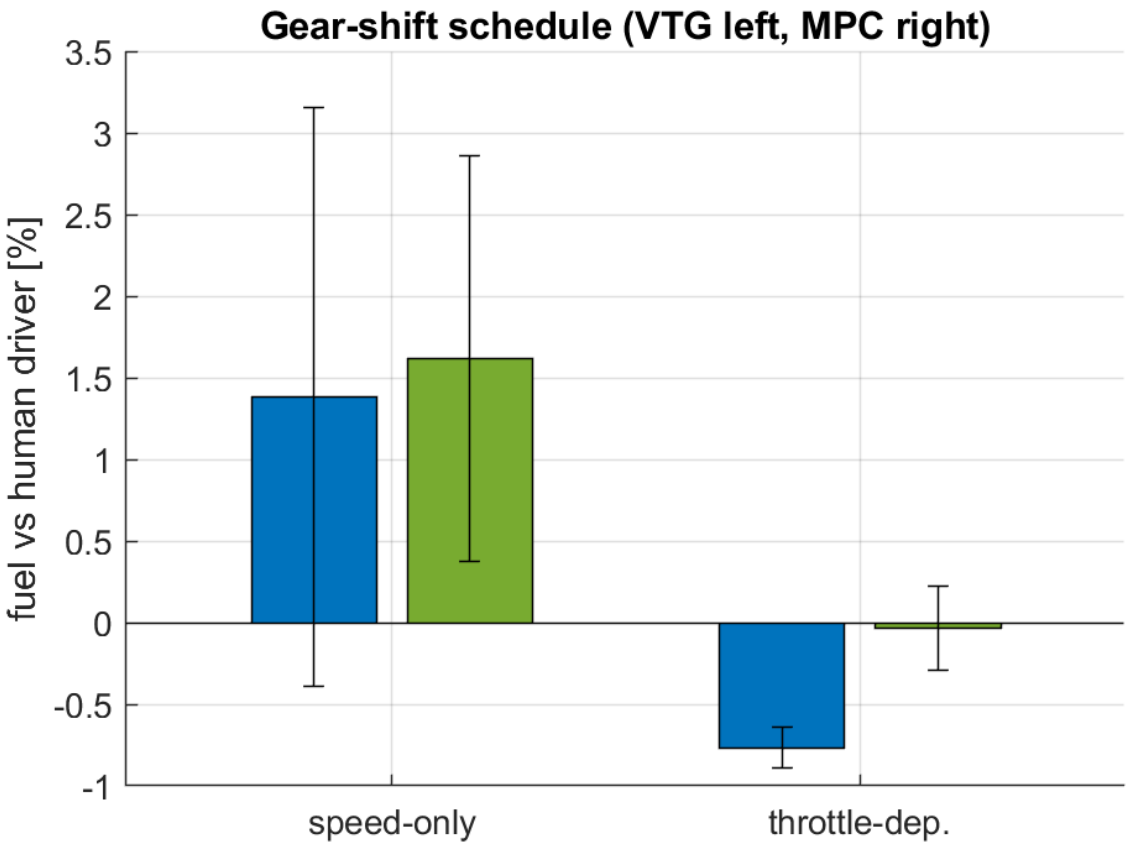


Figure 5.3: Cruise engine speed at equal vehicle speed and resulting fuel comparison under the two gear-shift schedules.

work, and quantifying it is listed among the future developments.

5.2 Case 2: Emergency Takeover and Safety Intervention Results

The Case 2 campaign evaluates three configurations over the same five seeds: a distracted driver without assistance, which is the reference case, a takeover to VTG and a takeover to MPC. All quantities in this section are evaluated after the distraction onset t_{switch}; the whole-run indicators are not reported, because up to the takeover all three configurations are driven by the same human model and are identical by construction. Table 5.5 collects the results.

5.2.1 Takeover Safety and Collision Avoidance

The distracted driver collides with the leading vehicle in **every one of the five seeds**. The collision occurs on average 62.3±29.1 s after the onset of the distraction: in four seeds at the anchored emergency deceleration, and in one seed against a leader decelerating

Table 5.5: Case 2 results after the distraction onset. Mean ± standard deviation over five paired seeds. For the unassisted runs the distracted phase extends to the collision, so the indicators before the takeover and those after the onset describe the same window.

Indicator	**No assistance**	**VTG**	**MPC**
Collisions	**5/5**	**0/5**	**0/5**
Time to collision event (s)	62.3 ± 29.1	—	—
Takeover delay (s)	—	48.3 ± 33.6	48.3 ± 33.6
Min TTC before takeover (s)	0.013 ± 0.002	2.28 ± 0.26	2.28 ± 0.26
Min gap before takeover (m)	0.09 ± 0.06	11.9 ± 5.5	11.9 ± 5.5
Min TTC after takeover (s)	—	$\mathbf{2.09 \pm 0.42}$	1.65 ± 0.69
Harshest braking after takeover (m/s^2)	—	$\mathbf{-6.05 \pm 1.47}$	-6.89 ± 0.43
Recovery gap (m)	—	14.2 ± 2.7	14.2 ± 2.7
Min gap after onset (m)	0.09 ± 0.06	3.11 ± 0.16	3.21 ± 0.23
Min TTC after onset (s)	0.003 ± 0.002	$\mathbf{2.09 \pm 0.42}$	1.59 ± 0.62
Consumption after onset, ICE (L/100 km)	—	**5.37**	5.40
Consumption after onset, EV (kWh/100 km)	—	**7.32**	7.36

towards a traffic-light stop, well before its hazard. This is the expected consequence of a 2.0 s reaction delay in dense urban traffic, and it establishes the reference against which the intervention is measured: without assistance, the outcome is not a degraded safety margin, it is contact.

With the supervisor active, **no collision occurs in any seed, with either controller**. The takeover restores a minimum gap of 3.11 ± 0.16 m (VTG) and 3.21 ± 0.23 m (MPC) and a minimum time-to-collision above 1.5 s: the emergency recovery restores margins of the same order as those the ADAS controllers hold in continuous operation, and well above those of the attentive human driver of Case 1.

Two properties of the campaign confirm that the comparison is clean. First, the takeover delay and all pre-takeover indicators are identical for the two controllers (48.3 ± 33.6 s, TTC $= 2.28 \pm 0.26$ s, gap $= 11.9 \pm 5.5$ m), because the human drives until the engagement and the supervisor decision depends only on the sensed state, not on the controller that will be engaged. Second, the leader trajectory is identical across the three configurations within each seed.

The takeover delay deserves attention. Averaged over seeds it is 48 s, but the quantity that matters is the margin between the engagement and the instant at which the reference run collides. In three of the five seeds that margin is approximately 1.6 s: a purely risk-triggered supervisor, which observes the danger and not the driver, intervenes at nearly the last instant at which the collision is still avoidable. The intervention is successful in all seeds, but with little reserve, and this quantifies the safety value of driver-state monitoring: a supervisor that detected the distraction itself would engage tens of seconds earlier, and the recovery would not require braking close to the friction limit.

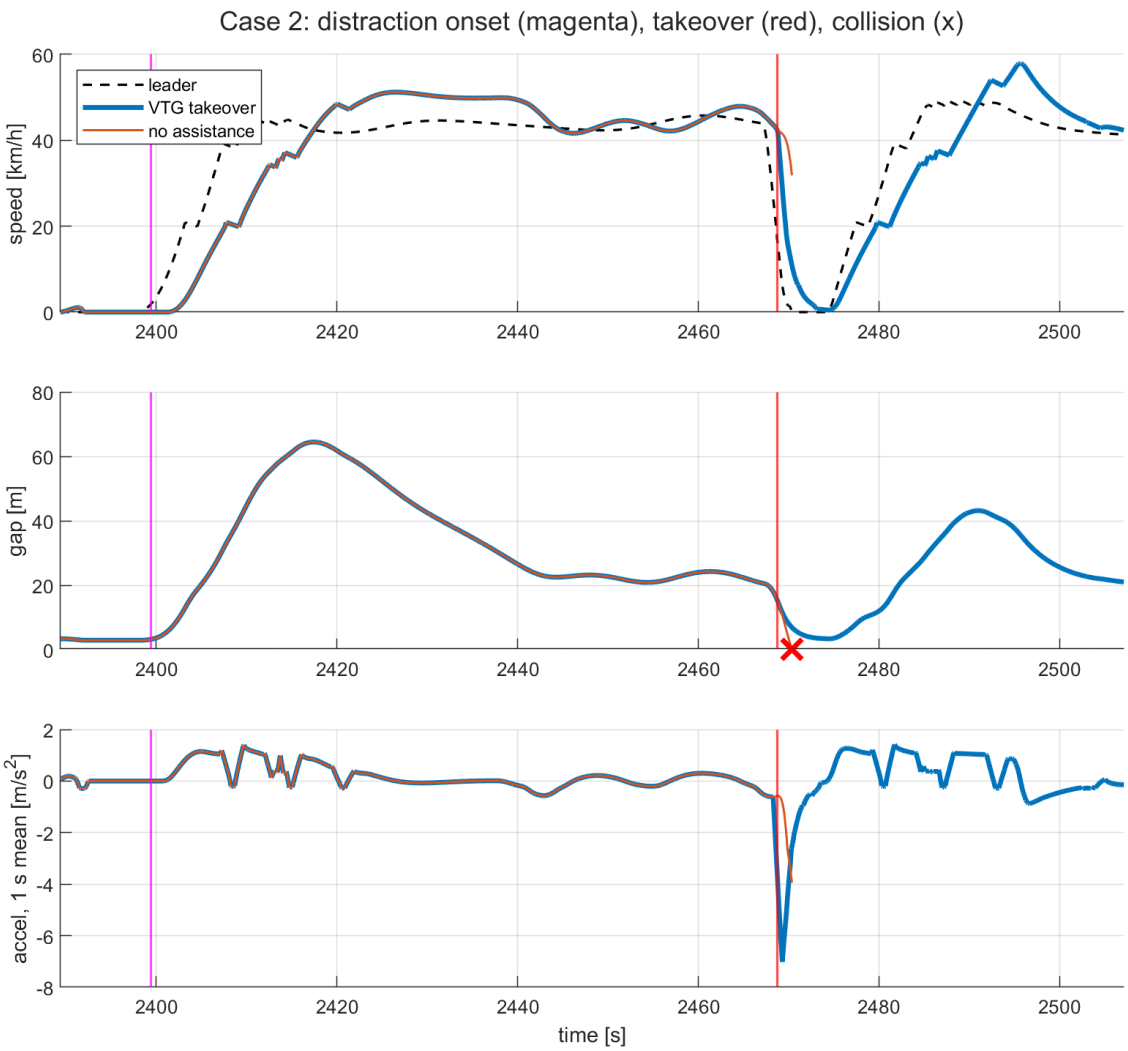


Figure 5.4: One seed of Case 2: the distracted driver without assistance closes the gap to contact, while the same run with the risk-triggered takeover recovers. Magenta marks the distraction onset, red the engagement of the ADAS, and the cross the collision. The two runs are identical until the engagement, so the unassisted trace is visible only over the last seconds.

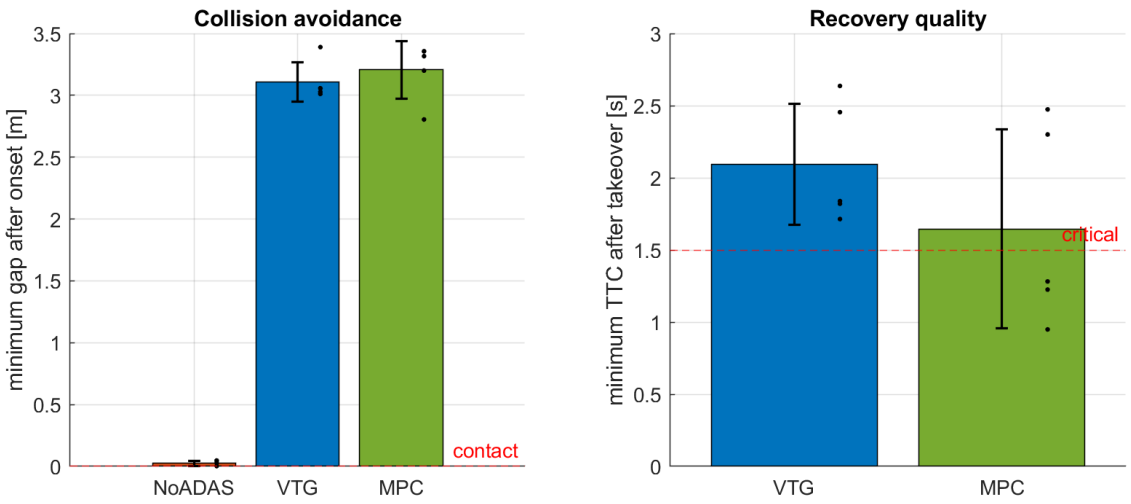


Figure 5.5: Case 2 outcomes over the five seeds: minimum gap after the distraction onset, where the unassisted runs reach contact, and minimum time-to-collision after the takeover. Dots are individual seeds.

5.2.2 Handover Transient and Controller Comparison

After the engagement both controllers execute an emergency deceleration close to the friction limit, then reopen the gap to 14.2 ± 2.7 m before settling on their spacing policy. The two controllers differ systematically in the quality of this recovery. The VTG law preserves 2.09 ± 0.42 s of minimum time-to-collision against 1.65 ± 0.69 s for the MPC; the paired difference is 0.45 ± 0.27 s in favour of VTG and has the same sign in all five seeds (paired t-test, $t = 3.65$, $p = 0.022$). The VTG law achieves this while braking *less* hard (-6.05 against -6.89 m/s^2), and its worst case over the campaign leaves 1.71 s of time-to-collision, whereas the worst MPC case leaves 0.95 s and a minimum gap of 2.80 m.

The ordering therefore reproduces, under an emergency transient, the ordering already observed in Case 1 under steady car-following. The reason is the same: the MPC reacts to a filtered estimate of the leader acceleration over a four-second linear prediction, while the VTG law acts algebraically on the instantaneous spacing and relative speed. When the situation develops in one or two seconds, the prediction adds latency rather than anticipation.

Finally, the consumption measured over the phase following the distraction onset is 5.37 against 5.40 L/100 km on the ICE powertrain and 7.32 against 7.36 kWh/100 km on the EV powertrain, that is 0.5 to 0.6% in favour of VTG, in the same direction and of the same magnitude as in Case 1. The consumption of the unassisted runs is not comparable, since those runs end prematurely at the collision and cover only part of the route.

As in Case 1, the ICE and EV campaigns produce identical trajectories seed by seed, so every safety indicator of Table 5.5 is common to both powertrains and only the energy rows differ. This was verified run by run and is a consistency check on the whole simulation

campaign.

Chapter 6

Conclusion

This thesis has presented the design, implementation and evaluation of a multi-objective Adaptive Cruise Control framework on a nonlinear 4-DOF longitudinal vehicle model in MATLAB/Simulink, comparing a reaction-delayed human driver model against a Variable Time Gap controller and a linear Model Predictive Controller over a stochastic urban driving cycle and two powertrain energy paths.

6.1 Concluding Remarks

1. **The benefit of longitudinal automation is primarily a safety benefit.** Over five paired seeds, the minimum time-to-collision increased from 1.16 ± 0.05 s under manual control to 2.90 ± 0.01 s with VTG and 2.20 ± 0.08 s with MPC, while the harshest sustained deceleration fell from -7.27 ± 0.19 m/s^2 to -4.06 ± 0.36 and -5.02 ± 0.87 m/s^2 respectively, and the time spent braking harder than -3.5 m/s^2 was approximately halved. These gains were obtained at an equal or slightly tighter median time gap and at equal distance covered, so they are not paid for in throughput.

2. **Disturbance attenuation is the underlying mechanism.** The manual driver amplifies the leader deceleration by about 20% because of its one-second reaction delay, and its positive kinetic energy exceeds that of the vehicle it follows. The VTG controller reproduces or reduces the leader positive kinetic energy, and therefore behaves as a string-stable element on a physically detailed plant.

3. **Energy savings are real but of the order of one percent, not tens of percent.** With respect to the human driver, VTG saves $1.69 \pm 0.17\%$ on the EV powertrain and $0.77 \pm 0.12\%$ on the ICE powertrain; the corresponding MPC figures are $0.61 \pm 0.35\%$ and a value indistinguishable from zero. Earlier versions of this study reported a fuel penalty of manual driving close to 48%; that figure was an artefact of a poorly tuned oscillating regulator used as the human baseline, and it does not survive the replacement of that baseline with a calibrated car-following model.

4. **On a conventional powertrain, the energy outcome is governed by the engine operating point, not by the trajectory.** The phase decomposition shows that smooth automated driving trades efficient acceleration time for inefficient light-load cruise time, and the gear-shift sensitivity study shows that the sign of the ICE comparison changes with the shift schedule: under a speed-only schedule the ADAS controllers consumed 1.4% to 1.6% more fuel than the human driver, and under a throttle-dependent schedule they consume the same or less. Any claim about the fuel benefit of an ACC strategy on a conventional vehicle must therefore state the transmission control strategy under which it was obtained.

5. **The simpler controller is sufficient for this task.** The VTG law outperformed the linear MPC on safety, comfort, repeatability and energy, while requiring only an algebraic evaluation per step instead of a constrained quadratic program solved at 10 Hz. The predictive controller pays for a four-second linear prediction on a cycle whose critical events are not predictable within that horizon.

6. **Automated intervention converts a certain collision into a safe recovery.** In Case 2 the distracted driver, with a reaction delay of 2.0 s, collided in all five seeds, on average 62 s after the onset of the distraction. With a risk-triggered takeover no collision occurred in any seed, with either controller, and the platoon was restored to a minimum gap above 3 m. The same controller ordering appeared in the recovery transient as in continuous operation: VTG retained 2.09 ± 0.42 s of time-to-collision against 1.65 ± 0.69 s for the MPC ($p = 0.022$), while braking less hard. The intervention, however, succeeded with little reserve: in three of five seeds the supervisor engaged approximately 1.6 s before the instant of the reference collision, because a risk-triggered policy observes the danger rather than the driver.

6.2 Limitations

The following limitations bound the validity of these results and should be taken into account when interpreting them.

- **Statistical scope.** Five seeds are sufficient to establish the safety results, whose effect sizes are an order of magnitude larger than their dispersion, and the EV energy differences, which are consistent in sign across all seeds. They remain a small sample for the ICE comparison.

- **Human model.** The IIDM parameters are standard literature values rather than a calibration against measured trajectories of drivers in Torino, and a single reaction delay of 1.0 s represents the attentive driver. Since the safety result is driven by this delay, its sensitivity is the first item of future work.

- **Engine braking.** The plant produces no drag torque at closed throttle, so gentle decelerations are obtained with a small positive throttle rather than by coasting in fuel cut. Approximately 15% of the fuel is consumed in this state. This affects all controllers similarly, but it compresses the difference between them.

- **Idle behaviour.** At standstill the engine speed decays to zero instead of settling at an idle speed; the fuel model compensates by charging a constant idle flow below 1 m/s, so the consumption figures are unaffected, but the modelled engine state at standstill is not physical.
- **EV model.** Constant drive and regeneration efficiencies, no battery mass, no thermal effects and no speed-dependent machine efficiency; the EV results quantify the influence of the driving style on an electrified energy path rather than the consumption of a specific production vehicle.
- **Scenario scope.** A two-vehicle platoon, longitudinal dynamics only, no cut-ins, no road grade, perfect sensing with no measurement noise or latency other than the modelled actuation lag.

6.3 Future Work

- **Driver-state monitoring.** Repeating the Case 2 campaign with a supervisor that engages on the *detected distraction* rather than on the developing risk would quantify the safety reserve gained by driver monitoring. Given that the risk-triggered policy engages roughly 1.6 s before the reference collision, this is expected to be the largest single safety improvement available within the present architecture.
- **Reaction-time sensitivity.** Repeating the Case 1 campaign with $\tau_r \in \{0.7, 1.0, 1.5\}$ s, and the Case 2 campaign with a distracted delay of 1.5 and 2.5 s, would quantify how much of the safety benefit of automation is attributable to the delay of the human baseline, and would place both cases on a single continuous axis.
- **Headway policy sweep.** Comparing the controllers at matched time gaps between 1.0 and 3.0 s would separate the contribution of the headway policy from that of the control law, and would quantify the throughput cost of safety.
- **Plant validation.** Standard manoeuvres – 0 to 100 km/h acceleration, coast-down and full braking – compared against published data for a vehicle of this class would strengthen the credibility of the absolute energy figures.
- **Computational cost.** Measuring the solver time of the MPC through an offline replay would complete the complexity argument of Chapter 5 with quantitative evidence.
- **Platoon extension.** Extending the simulation to three or more vehicles would allow the string stability property suggested by the positive kinetic energy comparison to be verified directly through the propagation of a disturbance.
- **Transmission co-design.** Since the ICE result is governed by the shift schedule, co-designing the ACC controller and the shift map – for example by allowing the controller to request an upshift when it plans a cruise phase – is a promising direction for genuine eco-driving gains on conventional powertrains.

Ultimately, this study shows that longitudinal automation on an urban cycle eliminates the critical collision risk introduced by human reaction delay and yields a small but consistent energy benefit, and that the magnitude of that benefit is determined as much by the powertrain and transmission models used to evaluate it as by the control strategy itself.

Bibliography

[1] Diana Yanakiev and Ioannis Kanellakopoulos. Nonlinear spacing policies for automated heavy-duty vehicles. *IEEE Transactions on Vehicular Technology*, 47(4):1365–1377, 1998.

[2] Martin Treiber, Arne Kesting, and Dirk Helbing. Delays, inaccuracies and anticipation in microscopic traffic models. *Physica A: Statistical Mechanics and its Applications*, 360(1):71–88, 2006.

[3] Chaofeng Pan, Aibao Huang, Jian Wang, Liao Chen, Jun Liang, Weiqi Zhou, Limei Wang, and Jufeng Yang. Energy-optimal adaptive cruise control strategy for electric vehicles based on model predictive control. *Energy*, page 122793, 2022.

[4] Mattia Boggio, Carlo Novara, and Michele Taragna. Trajectory planning and control for autonomous vehicles: a "fast" data-aided nmpc approach. *European Journal of Control*, page 100857, 2023.

[5] Sahand Mosharafian and Javad Mohammadpour Velni. A hybrid stochastic model predictive design approach for cooperative adaptive cruise control in connected vehicle applications. *Control Engineering Practice*, page 105383, 2023.

[6] Ronghui Zhang, Na Wu, Zihan Wang, Kening Li, Zhumei Song, Zhenting Chang, Xia Chen, and Fan Yu. Constrained hybrid optimal model predictive control for intelligent electric vehicle adaptive cruise using energy storage management strategy. *Journal of Energy Storage*, page 107383, 2023.

[7] Zahra Mehraban, Ashkan Y Zadeh, Hamid Khayyam, Rammohan Mallipeddi, and Ali Jamali. Fuzzy adaptive cruise control with model predictive control responding to dynamic traffic conditions for automated driving. *Engineering Applications of Artificial Intelligence*, page 109008, 2024.

[8] Shiying Dong, Xi Luo, Yuxiang Zhang, Qifang Liu, Bingzhao Gao, and Hong Chen. Computationally efficient predictive cruise control of electric vehicles with nonsmooth system. *Control Engineering Practice*, page 106466, 2025.

[9] Shaimaa K. El-Baklish, Anastasios Kouvelas, and Michail A. Makridis. Driving towards stability and efficiency: A variable time gap strategy for adaptive cruise control. *Transportation Research Part C: Emerging Technologies*, page 105074, 2025.

[10] Yuxi Li, Yu Zhang, and Gang Hao. Stochastic energy model predictive control for cooperative adaptive cruising. *Energy*, page 137573, 2025.

[11] Sahand Mosharafian and Javad Mohammadpour Velni. Event-triggered robust model predictive design for cooperative adaptive cruise control application. *IFAC PapersOnLine*, pages 253–258, 2025.

[12] Zhou Wang, Ronghui Zhang, Shuai ... Niu, Dong Zhang, and Bolin Gao. Cloud-based energy-efficient predictive cruise control accounting for nonlinear vehicle dynamics on highways. *Energy*, page 139984, 2026.

[13] Giancarlo Genta and Lorenzo Morello. *The Automotive Chassis, Volume 2: System Design*. Springer, Dordrecht, 2009.

[14] Rajesh Rajamani. *Vehicle Dynamics and Control*. Springer, New York, 2 edition, 2012.

[15] Martin Treiber and Arne Kesting. *Traffic Flow Dynamics: Data, Models and Simulation*. Springer, Berlin, Heidelberg, 2013.