



UNIVERSITY OF LEEDS

This is a repository copy of *Comparing and validating models of driver steering behaviour in collision avoidance and vehicle stabilisation*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/102231/>

Version: Accepted Version

Article:

Markkula, G orcid.org/0000-0003-0244-1582, Benderius, O and Wahde, M (2014) Comparing and validating models of driver steering behaviour in collision avoidance and vehicle stabilisation. *Vehicle System Dynamics*, 52 (12). 12. pp. 1658-1680. ISSN 0042-3114

<https://doi.org/10.1080/00423114.2014.954589>

© 2014 Taylor & Francis. This is an Accepted Manuscript of an article published by Taylor & Francis in *Vehicle System Dynamics* on 12 September 2014, available online: <http://www.tandfonline.com/10.1080/00423114.2014.954589>. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Unless indicated otherwise, fulltext items are protected by copyright with all rights reserved. The copyright exception in section 29 of the Copyright, Designs and Patents Act 1988 allows the making of a single copy solely for the purpose of non-commercial research or private study within the limits of fair dealing. The publisher or other rights-holder may allow further reproduction and re-use of this version - refer to the White Rose Research Online record for this item. Where records identify the publisher as the copyright holder, users can verify any specific terms of use on the publisher's website.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

(Final draft post-refereeing, pre-publication)

RESEARCH ARTICLE

Comparing and validating models of driver steering behaviour in collision avoidance and vehicle stabilization

G. Markkula^{a,b,*}, O. Benderius^b, and M. Wahde^b

^a*Transport Analysis, Volvo Group Trucks Technology, Advanced Technology and Research, Göteborg, Sweden;* ^b*Department of Applied Mechanics, Chalmers University of Technology, Göteborg, Sweden*

(Received Month Day Year)

A number of driver models were fitted to a large data set of human truck driving, from a simulated near-crash, low-friction scenario, yielding two main insights: Steering to avoid a collision was best described as an open-loop manoeuvre of predetermined duration, but with situation-adapted amplitude, and subsequent vehicle stabilization could to a large extent be accounted for by a simple yaw rate nulling control law. These two phenomena, which could be hypothesized to generalize to passenger car driving, were found to determine the ability of four driver models adopted from literature to fit the human data. Based on the obtained results, it is argued that the concept of internal vehicle models may be less valuable when modelling driver behaviour in non-routine situations such as near-crashes, where behaviour may be better described as direct responses to salient perceptual cues. Some methodological issues in comparing and validating driver models are also discussed.

Keywords: Driver models, steering, collision avoidance, low friction, driving experience, electronic stability control

Manuscript word count: 7501 words including appendix

1. Introduction

It is well established that driver behaviour plays a prominent role in the causation of traffic accidents [1, 2], and considerable research effort has been spent on understanding and describing driver behaviour in near-crash situations. This is not an easy object of study, but as a result of accident reconstructions, large-scale naturalistic data collection projects, and experiments on test tracks and in driving simulators, there is a growing body of knowledge on the various reasons why drivers end up in critical situations, such as inattention [3] or incorrect expectations [4, 5], and on how drivers typically control the vehicle if and when they try to avoid an imminent crash [6–9].

An important application of such knowledge is the construction of quantitative models of driver control behaviour in near-crash situations. When put to use in computer simulations, such models permit cost-efficient safety performance optimization of, for example, infrastructure designs [10], vehicle designs [11], or active support systems that provide warnings or control interventions [12–14].

*Corresponding author. E-mail: gustav.markkula@volvo.com

There is a wealth of existing driver behaviour models that could be useful in such simulation-based research efforts, reviewed in [15–19]. However, a recent review, focusing specifically on models that have been applied in simulation of near-collision situations [20], noted two clear limitations in the literature: (a) With just a few exceptions [21–23], new models have been proposed without comparing their behaviour to that of existing, alternative models, making it difficult to know which models to prefer for a given application. (b) Validation of model behaviour against human behaviour data from real or reasonably realistic near-crash situations has been virtually non-existent. Many models of steering control were found to have been validated against human behaviour in predefined test-track manoeuvres, for example a double lane change [24]. However, such tests seem rather unlike real-life, unexpected near-crash situations, and could potentially elicit qualitatively different behaviours from drivers [25–27].

This paper addresses both of the two above-mentioned limitations, in the specific context of collision avoidance and subsequent vehicle stabilization, on a low friction road surface. One use for models validated in this type of context is simulation-based evaluation of vehicle stability support systems such as *electronic stability control* (ESC) [28–31]. In [32], it was shown that one existing driver model could reproduce the stabilization steering behaviour observed after unexpected and expected near-collisions in a driving simulator study, previously described in [33]. Here, the analysis of this dataset, fitting models separately to each human driver, will be extended to include also the collision avoidance phase of the studied scenario, and to include a comparison of a number of existing and novel models of steering.

The next section will describe the data collection simulator study and the driver models, as well as the method for fitting the models to the human steering data. Then, model-fitting results will be provided, including some analysis of the obtained model parameters. The subsequent discussion will highlight differences between the models and their respective strengths and weaknesses in the studied scenario, as well as some challenges involved in model comparison and validation.

2. Method

2.1. Data collection

The human driving data used here were collected in the moving-base driving simulator *VTI Simulator II* in Linköping, Sweden. Full details on the simulator and the experimental procedures adopted in this study can be found in [33]. In summary, 48 drivers, driving a three-axle rigid simulated truck (6.2 m from first to last axle) at 80 km/h, experienced an unexpected lead vehicle deceleration scenario on a low-friction ($\mu = 0.25$) road surface. Half of the subject drivers subsequently also experienced the same scenario an additional twelve times each, in a novel paradigm for repeated collision avoidance, and it is this 24-driver repeated-scenario data set which is used here. Half of the 24 drivers were novices, who had just obtained, or were just about to obtain, their heavy truck driving license, and half were experienced drivers, with at least six years of professional experience in commercial operations. In half of all measurements, the simulated truck had an active ESC system, a software-in-the-loop implementation of the actual Volvo Trucks on-market ESC, and in the other half of measurements drivers were aided only by the anti-lock braking system (ABS). The critical scenario, requiring a steering manoeuvre for successful collision avoidance, is illustrated in Fig. 1, together with an overview of the observed vehicle trajectories.

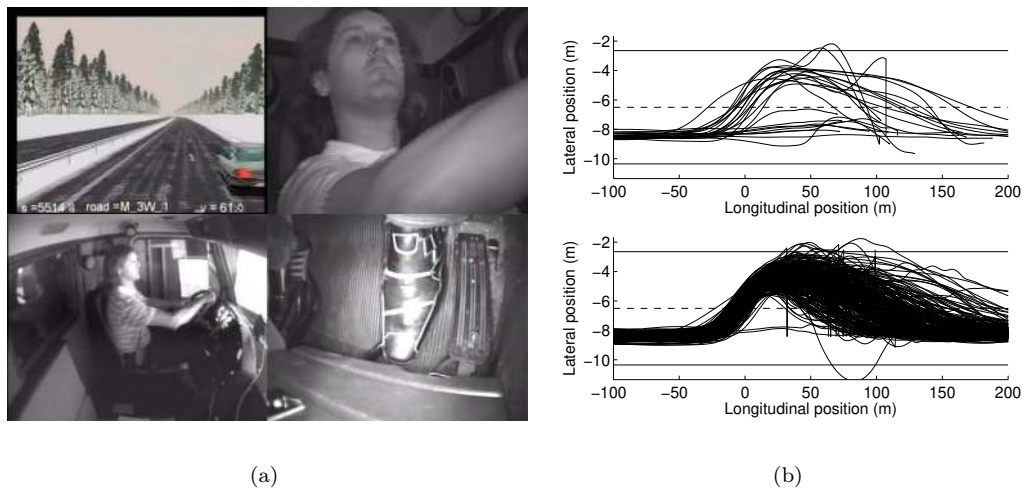


Figure 1. The studied critical scenario. (a) A screenshot from the experiment video log, showing the two-lane winter road on which the maneuvering took place, the braking lead vehicle, and a subject engaged in steering collision avoidance. (b) The totality of observed vehicle trajectories, in the unexpected (top panel) and repeated (bottom panel) scenarios. Horizontal gray lines show lane boundaries, and each black line shows the movement of the front center of the truck in one recorded event. Longitudinal position zero corresponds to the point where the front of the truck reached the rear of the lead vehicle. The unexpected scenario data are not used in the model-fitting analyses presented in this paper, but are shown here to illustrate that the two scenarios generated roughly similar human behaviour; see further [32] and [34].

It has been demonstrated elsewhere that the unexpected and repeated scenarios generated similar initial steering avoidance situations [33, 34], and elicited similar driving steering behaviour, both during collision avoidance [34] and stabilization [32].

2.2. Tested models

The set of driver models to test was defined so as to include both some well-known path-following models of steering, often available in off-the-shelf software for e.g. simulating predefined manoeuvres (the MacAdam and Sharp *et al.* models), as well as models of routine lane-keeping which may be less familiar but which take different, and in our view promising, modelling approaches (the Salvucci & Gray and Gordon & Magnuski models). Furthermore, based on the results from parameter-fitting these existing models to the behaviour of the human drivers, two very simple additional models of steering were developed, one targeting collision avoidance only, and the other targeting only vehicle stabilization. Below, all tested models will be briefly described, along with specific implementation details when needed. For ease of reading, consistent notation is used for quantities that are shared across models, in some cases departing from the symbols used by the original authors.

2.2.1. The MacAdam model

At a given time t , the model proposed by MacAdam [35], illustrated in Fig. 3(a), applies the steering wheel angle $\delta(t)$ that minimizes the predicted lateral deviation from a *desired path*, by minimizing the following functional:

$$J(t) = \int_{t-T_R}^{t-T_R+T_P} (f(\eta) - y(\eta))^2 d\eta \quad (1)$$

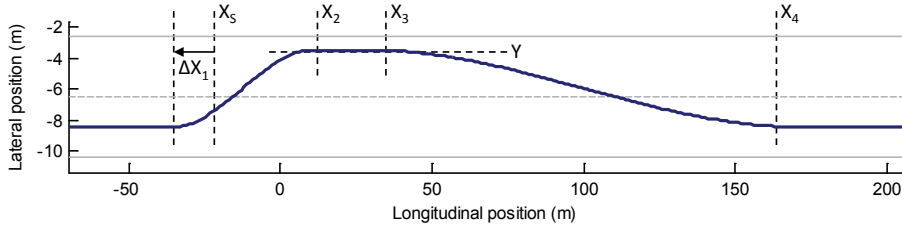


Figure 2. The desired path used for the MacAdam [35] and Sharp *et al.* [36] models, with ΔX_1 , X_2 , X_3 , X_4 , and Y as driver-specific model parameters. The two lane changes are cubic splines with lateral speed zero at beginning and end. Before and after the manoeuvre, desired lateral position is set to the middle of the right driving lane. Longitudinal position zero corresponds to the point where the front of the truck reached the rear of the lead vehicle, and X_5 is the longitudinal position at which the human driver's collision avoidance steering reached half of its maximum value, in a specific recorded instance of the critical scenario.

In Eq. (1), $y(t)$ and $f(t)$ are the predicted and desired lateral positions, and T_R and T_P are model parameters corresponding to a *reaction time* and a *preview time*, respectively. Here, the desired path $f(t)$ of the vehicle was defined as shown in Fig. 2, using model parameters ΔX_1 , X_2 , X_3 , X_4 , and Y . In order to allow for intra-driver variability in the exact point of collision avoidance initiation, the lane change to the left was set to start at a longitudinal position $X_5 - \Delta X_1$, where X_5 was the longitudinal position at which the steering wheel reached half of its maximum leftward deflection (i.e. X_5 had a unique value for every recorded instance of the critical scenario).

The MacAdam model's prediction $y(t)$ of lateral position relies on a linear *internal vehicle model*. Here, the same classical type of one-track model as used by MacAdam [35] was adopted, but with three axles instead of two:

$$\dot{\mathbf{x}} = \mathbf{F}\mathbf{x} + \mathbf{g}\delta = \begin{bmatrix} -\frac{C_{\alpha f} + C_{\alpha m} + C_{\alpha r}}{m v_x} & -\frac{a C_{\alpha f} + b C_{\alpha m} + c C_{\alpha r}}{m v_x} - v_x \\ -\frac{a C_{\alpha f} + b C_{\alpha m} + c C_{\alpha r}}{I_z v_x} & -\frac{a^2 C_{\alpha f} + b^2 C_{\alpha m} + c^2 C_{\alpha r}}{I_z v_x} \end{bmatrix} \mathbf{x} + \begin{bmatrix} \frac{G C_{\alpha f}}{m} \\ \frac{G a_f C_{\alpha f}}{I_z} \end{bmatrix} \delta \quad (2)$$

where $\mathbf{x} = [v_y \ \dot{\psi}]^T$, v_x and v_y are longitudinal and lateral speeds in the vehicle's reference frame, and ψ is the yaw angle of the vehicle. The three C_{α} parameters and a , b , c are tire cornering stiffnesses and longitudinal distances to the vehicle's mass centre, for the front, middle, and rear axles, respectively. The parameters m and I_z are vehicle mass and moment of inertia, and G is the steering gear ratio.

Since the linear vehicle model cannot account well for skidding, it was parameter-fitted only to recordings with maximum body slip angle $\beta < 1^\circ$ (3 % of the total data set). This can be understood as assuming that drivers had acquired an understanding of vehicle dynamics from normal, high-friction driving, and applied this understanding also during yaw instability¹. Only the three cornering stiffnesses were fitted to the data; the other parameters were taken from the non-linear model used in the simulator study. Fig. 4 illustrates the resulting model performance at various magnitudes of yaw instability.

2.2.2. The Sharp *et al.* model

The model proposed by Sharp *et al.* [36] also makes use of the desired path construct, but, as illustrated schematically in Fig. 3(b), instead calculates its steering wheel input as a weighted sum of current and previewed path deviations e_i along a forward *optical lever*, extending a preview time T_P ahead, and the current deviation

¹Various approaches were explored for fitting the linear model also to recordings with more severe yaw instability, but were not found to improve the fit of the resulting driver model.

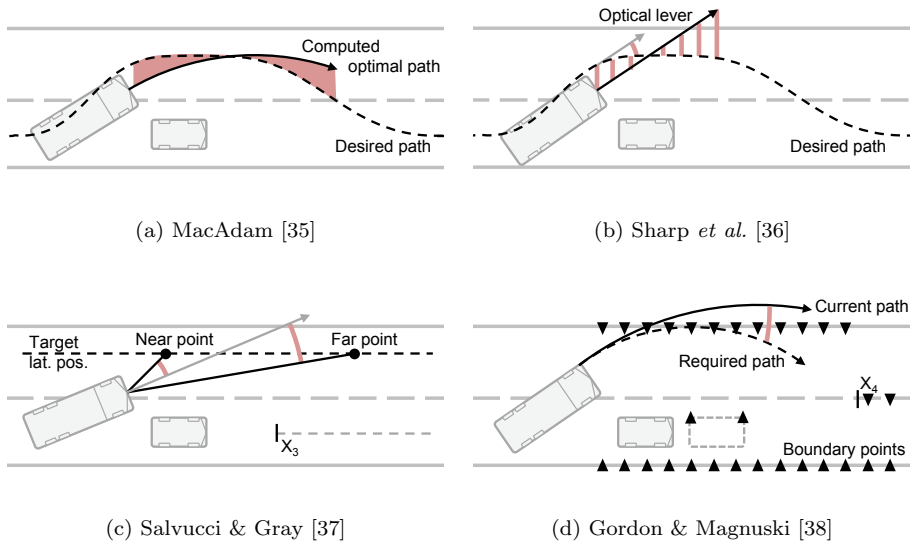


Figure 3. Schematic illustrations of the models adopted from literature.

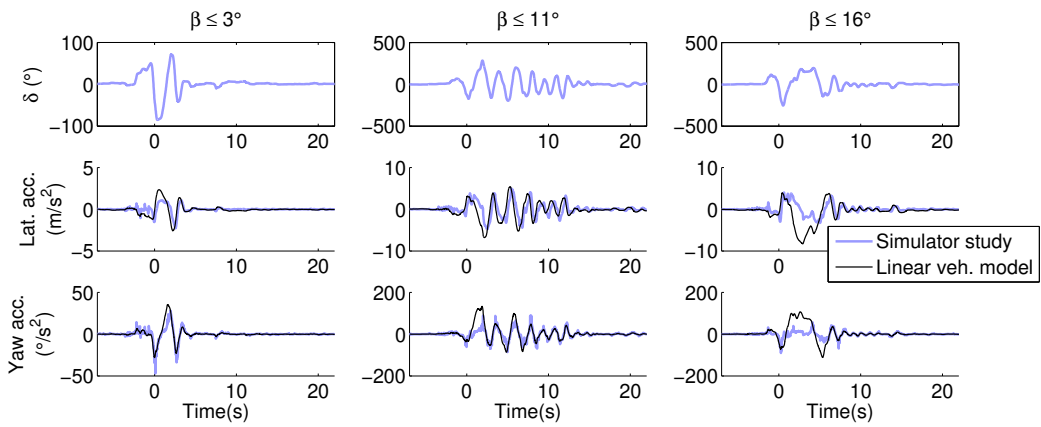


Figure 4. The top row of panels show observed steering wheel movements in three recordings of the repeated scenario, with maximum attained body slip angles β increasing from left to right. The bottom two rows of panels show, for the same recordings, the observed vehicle dynamics from the full non-linear vehicle dynamics model used in the data collection experiment, compared with the vehicle dynamics predicted by the linear model used with the MacAdam driver model in this paper.

e_ψ between vehicle and path heading:

$$\delta = K_\psi e_\psi + K_1 e_1 + K_p \sum_{i=2}^n K_i e_i \tag{3}$$

Here, K_ψ , K_1 , and K_p were treated as free model parameters, whereas the number n of preview points, their spacing along the optical lever, and the exponentially decreasing profile for the preview gains K_i (with $2 \leq i \leq n$) were adopted from [36]. Additionally, to allow for a fair comparison with the other models, a reaction time delay parameter T_R was added to Eq. (3). The saturation functions included by Sharp *et al.* with the purpose of “preventing the steer angle from exceeding a reasonable range” [36, p. 312], were not included.

2.2.3. The Salvucci and Gray model

The model by Salvucci & Gray [37] is mathematically rather similar to the Sharp *et al.* model. However, instead of being derived from linear optimal control theory, it builds on experiments and modelling in psychology, motivating: (a) the use of the rate of change $\dot{\delta}$ of steering as an input variable rather than δ [39], and (b) the separation of controlled quantities (the input to the driver) into one *near point* and one *far point* [40] on a target lateral position, as illustrated in Fig. 3(c). It is assumed that the driver aims to keep the the sight angles θ_n and θ_f to these points stationary, while at the same time attempting to reduce the near point angle to zero:

$$\dot{\delta} = k_{nP}\dot{\theta}_n + k_f\dot{\theta}_f + k_{nI}\theta_n \quad (4)$$

In addition to the gain parameters in Eq. (4), free parameters were also included for the longitudinal distances D_n and D_f to the near and far points, respectively¹, as well as a reaction time T_R . Analogously to the desired path of the models described above, the target lateral position was initially set to the middle of the right lane, then set to a position Y when the truck reached longitudinal position $X_S + \Delta X_I$, and then back to the middle of the right lane at longitudinal position X_3 . To test the sensitivity of the model to the preview distance parameters, an additional version of the model was tested, where these parameters were fixed, for all drivers, at the median values $D_n = 16$ m and $D_f = 123$ m, observed in the optimizations where these parameters were left free.

2.2.4. The Gordon and Magnuski model

Since the models described above all aim at reducing the deviation from a desired path or lateral position to zero, they could be referred to as *optimizing* models. In contrast, the model by Gordon & Magnuski [38], illustrated in Fig. 3(d), operates in what can be called a *satisficing* [41] manner: It assumes that the driver is content with staying inside a delimited region, modelled using *boundary points*. Specifically, the model compares the current yaw rate to the yaw rates needed to steer clear of each boundary point, identifies the point with the greatest mismatch, and applies a rate of steering aimed at reaching, within a time τ_s , the required yaw rate $\dot{\psi}_{\text{req}}$ for this point, assuming a simple vehicle model with wheel base L :

$$\dot{\delta} = -\frac{L}{G\tau_s v_x}(\dot{\psi} - \dot{\psi}_{\text{req}}) \quad (5)$$

Before computing $\dot{\psi}_{\text{req}}$, the model also applies a vehicle state prediction to counteract its own reaction time delay T_R .

The original publication [38] considered only lane keeping, but Chang [42] applied the same model to avoidance of static obstacles, with a safety margin ρ_C . In the present work, seemingly the first time the model is applied to obstacles and lane boundaries simultaneously, conflicts with lead vehicle boundary points were given priority over lane boundary conflicts, and a separate safety margin ρ_L for lane boundaries was added, with allowed negative values in the optimization, to account for the apparent acceptance of moderate lane excursions in some of the

¹Following [40], Salvucci & Gray [37] specified preview in terms of angles down from the horizon, which in practice amounts to the same as using a preview distance. Here, it was also attempted to make the preview speed-dependent, as in the MacAdam and Sharp *et al.* models, e.g. $D_n = T_n v_x$, but if anything this reduced the model's ability of fitting the human data.

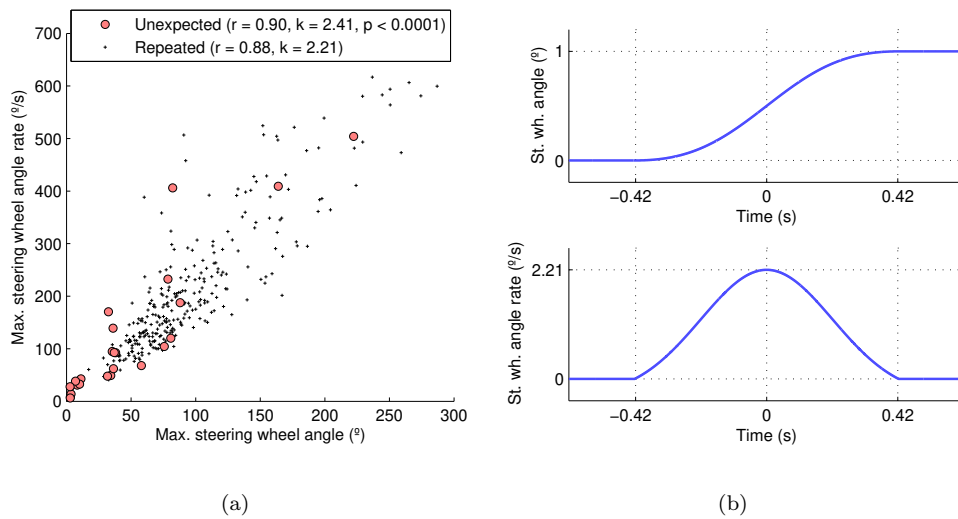


Figure 5. Open loop avoidance steering. (a) The observed correlation between maximum leftward steering wheel angle before reaching the lead vehicle, and maximum leftward steering wheel angle rate during the same period. The Pearson correlation coefficients r are provided, as well as the slope k of least-squares fit lines with zero intercept. The unexpected scenario data are not used in the model-fitting analyses presented in this paper, but are shown here to illustrate that the two scenarios generated roughly similar human behaviour; see further [32] and [34]. (b) The steering wheel input generated by the open loop avoidance model in this paper, here parameterized to illustrate the width of such a steering wheel pulse (0.84 s) that would yield the k observed for the repeated scenario in panel (a).

human drivers. Furthermore, extending the model to handle the non-static lead vehicle, lead vehicle state prediction was included, as also illustrated schematically in Fig. 3(d). In three different versions of the model, this prediction was done assuming a constant lead vehicle acceleration (2nd order), speed (1st order), and position (0th order), respectively. To implement the scenario studied here, the lead vehicle boundary points were included only from longitudinal position $X_S + \Delta X_I$, and the lane change back to the right lane was achieved by placing the left-side boundary points beyond longitudinal position X_4 to between the two driving lanes; again, see Fig. 3(d).

2.2.5. Open loop avoidance models

An additional model of collision avoidance steering was tested, motivated by a linear correlation previously reported by Breuer [25], and replicated here: As shown in Fig. 5(a), higher-amplitude avoidance manoeuvres were carried out with faster steering movements. This finding suggests (a) that avoidance manoeuvre duration was roughly constant between scenario recordings, and (b) that each manoeuvre’s amplitude was determined *before* its initiation.

Therefore, in contrast to the *closed-loop* models described above, which calculate a new control input at each time step in a simulation, an *open-loop* model was posited, applying a pulse of steering wheel rotation represented as a Gaussian cut off at ± 2 standard deviations; see Fig. 5(b). The pulse duration $T_D = 2T_H$ was included as a free parameter, and pulse amplitude was determined as a function of the collision situation a reaction time T_R before manoeuvre initiation. To allow for the above-mentioned intra-driver variation in collision avoidance timing, the peak of the pulse was placed at time $T_S + T_A$, where T_S was the time at which the truck’s longitudinal position was X_S (see above), and T_A was another free model parameter.

Five different versions of the model were tested, all using a model parameter K to determine the steering pulse amplitude as (a) a constant, situation-independent

amplitude K , (b) K times the *optical expansion*, or *looming*, of the lead vehicle on the driver's retina [43], or (c-e) K times the steering required to avoid the lead vehicle with a safety margin ρ_C , given the steady state yaw rate response of the linear vehicle model described in 2.2.1, and a lead vehicle state prediction of orders zero through two (cf. Sec. 2.2.4).

While this model is neutral on whether the driver is controlling steering wheel angle δ or its rate of change $\dot{\delta}$, the choice of target signal has an impact on model parameter-fitting (further discussed in Sect. 4.3). Therefore, this model was fitted separately both as controlling δ and $\dot{\delta}$.

The model's sensitivity to the T_R parameter was also tested, by parameter-fitting an additional version of the δ -controlling, 2nd order yaw rate requirement model, with T_R fixed at 0.2 s for all drivers.

2.2.6. Yaw angle/rate nulling stabilization models

As has been reported elsewhere [32], the Salvucci & Gray model is reasonably successful at fitting the stabilization steering data studied here. Additional exploration indicated that much of the variance explained by the model was accounted for by its far point control (the second term in Eq. 4), shown in [32] to approximate a *yaw rate nulling* steering behaviour: $\dot{\delta} = -K\psi$, where K is a model parameter. Here, such a model was tested directly, as well as a time-integrated *yaw angle nulling* version $\delta = -K\psi$, both with a reaction time delay T_R .

2.3. Division into avoidance and stabilization steering phases

Preliminary experimentation indicated that the steering models were differentially successful at fitting steering during collision avoidance and vehicle stabilization. Therefore, the data set was split accordingly, and model parameter-fitting was carried out separately on the two sets.

The *collision avoidance phase* of a recorded scenario was defined to begin when the lead vehicle started decelerating, and to end when the driver began applying considerable rightward steering wheel rotation, interpretable as a transition from leftward collision avoidance, to lane alignment and vehicle stabilization. This onset of rightward steering wheel rotation was generally clearly visible in the data, and was found to be suitably defined as the last point of leftward steering ($\delta > 0$) where $\dot{\delta} > -50^\circ/\text{s}$. The example recordings in Fig. 6 (further explained in Sect. 3) illustrate where this transition typically occurred.

The *stabilization phase* was defined to begin at the same transition point, and to end at whichever occurred first of (a) the truck having travelled 250 m after passing the lead vehicle, (b) the truck's longitudinal speed falling below 10 km/h, or (c) the driving simulator's safety shutdown system having aborted the scenario due to road departure, or a deviation of truck heading from the road's forward direction of 90° or more.

2.4. Model parameter-fitting

The repeated scenario generated 12 measurements for each of the 24 subject drivers except three, where, due to technical shortcomings, or subject failure to comply with experimental instructions, one or two scenario instances could not be recorded or used. Before parameter-fitting of models, all recorded signals were down-sampled to 5 Hz. The main motivation for down-sampling was the increase in optimization speed, but it could also be argued that at high sample rates, adjacent data points would anyway be highly correlated, and the input quantities to the tested models,

all constrained by the dynamics of the truck on the road, would also not be expected to contain much valuable information frequencies above 5 Hz.

At each included data point i , a model undergoing parameter-fitting was fed its required input data from the appropriate recorded signals, with delays if applicable, and the goal of the parameter-fitting, implemented using a *genetic algorithm* (GA), was to achieve a model output \hat{x}_i as close as possible to the observed human control x_i at the same point in time, with x_i equal to either δ_i or $\dot{\delta}_i$, depending on the model¹. For further information on GA optimization in general, see [44], and see Appendix A for full details on the specific GA used here.

To quantify model fit, the *coefficient of determination* R^2 , interpretable as the fraction of steering variance being explained by the model [45], was calculated for each scenario S as:

$$R^2 = 1 - \frac{SS_R}{SS_T} = 1 - \frac{\sum_{i \in S} (\hat{x}_i - x_i)^2}{\sum_{i \in S} (x_i - \bar{x}_S)^2} \quad (6)$$

where \bar{x}_S is the average of x_i , for $i \in S$. In other words, R^2 can be negative, if a model provides a worse model fit than simply guessing that $\hat{x}_i = \bar{x}_S$ for all $i \in S$.

Holdout validation [44] was adopted: For each subject driver, the set of available recorded scenarios was divided into one *training set* and one *validation set*, of equal size². The GA was set up to maximize the average of R^2 across the scenarios in the training set, but, in order to prevent *over-fitting*, the final model parameterization was selected as the parameterization with highest average R^2 across the validation set. The allocation of recorded scenarios to the two sets was designed to balance the amount of occurring vehicle instability between them: For each driver, the recorded scenarios were ordered by increasing maximum body slip angle, this ordered list was separated into pairs, and finally one randomly selected scenario in each pair was assigned to the training set, and the other to the validation set.

Thus, for each of the 24 drivers and each of the two steering phases, one optimization was carried out of each tested driver model. The total number of training and validation data points x_i used for one optimization ranged from 200 (avoidance steering for a subject with only ten recorded scenarios) to 900 (stabilization steering for a subject with twelve recorded scenarios).

3. Results

Table 1 summarizes the model-fitting results, per model and steering phase, as the average validation R^2 across the 24 drivers. Also listed are the numbers of effective free model parameters (N_{eff}), based on whether parameters were considered to have an effect on steering in the two phases; see Appendix A for full details. Note that the open loop avoidance model appears in the table both as a δ and $\dot{\delta}$ controlling model (see Sect. 2.2.5).

Figs. 6 and 7 show, in their leftmost columns, distributions of per-driver validation R^2 for some of the best-fitting model variants. It can be observed that the

¹Alternatively, one could have rerun the studied scenario in closed-loop simulation from initial conditions, fitting parameters to achieve a match between resulting driver steering histories or vehicle trajectories. Such an approach was not adopted here, both due to it being several orders of magnitude more computation-intensive, and since the inherent instability of the low-friction scenario would presumably have rendered fitting very difficult; a small error in driver model or initial conditions can lead to large deviations in scenario outcome.

²Except for one single subject driver where the number of available instances was odd, for which one more instance was allocated to the training set.

Table 1. Effective number of model parameters (N_{eff}) and average goodness-of-fit (R^2) on validation data, across all drivers, for all tested models in the avoidance and stabilization steering phases. Note that some models were tested in only one of the two phases.

Target signal	Model	Variant	Avoidance		Stabilization		
			N_{eff}	Av. R^2	N_{eff}	Av. R^2	
Steering wheel angle	MacAdam		6	0.49	5	0.50	
	Sharp et al.		9	0.46	8	0.60	
	Open loop avoidance	constant		3	-1.20		
		looming		4	0.71		
		0 th order		5	0.54		
		1 st order		5	0.66		
		2 nd order		5	0.76		
		2 nd order, T_R fixed		4	0.75		
	Yaw angle nulling				2	0.35	
	Steering wheel angle rate	Salvucci & Gray	preview free	8	0.20	8	0.68
preview fixed					6	0.68	
Gordon & Magnuski		0 th order	5	0.25	4	0.49	
		1 st order	5	-0.02	4	0.50	
		2 nd order	5	0.17	4	0.50	
Open loop avoidance		constant	3	0.41			
		looming	4	0.46			
		0 th order	5	0.47			
		1 st order	5	0.40			
		2 nd order	5	0.47			
Yaw rate nulling					2	0.54	

spread across drivers was rather similar between models. Further divisions into subgroups based on driver experience and ESC state indicated limited or no impact of these factors on model fit; one example of such a division can be seen in Fig. 9(a). Based on these observations, the discussion in the next section will compare models mainly in terms of the average validation R^2 values in Table 1. As a complement to this perspective, and to provide a more thorough grasp of actual model behaviour, Figs. 6 and 7 also show five example scenario recordings each, along with model predictions. These examples were selected to include driving both with and without ESC for both low and high experience drivers, to illustrate some specific strengths and weaknesses of the different models, while at the same time aiming for an average R^2 across the examples close to the average validation R^2 for each model. The discussion in the next section will provide suggestions on how to interpret the various examples.

Fig. 8 shows distributions of obtained parameters for the parameter-reduced variants of the open loop avoidance and Salvucci & Gray models, as well as for the yaw rate nulling model. The correlation between the safety margin ρ_C and the steering gain K in panel (a) is statistically significant¹ ($r = -0.76$; $p < 0.0001$), but the difference in the steering pulse duration T_H between experience groups in panel (b) is not ($t(22) = 0.654$; $p = 0.51$). The correlation between the steering gains k_f and k_{nP} in panel (c) is statistically significant ($r = -0.71$; $p = 0.0001$), and so is the difference in reaction time T_R between experience groups for the Salvucci & Gray model (panel (d); $t(22) = -2.19$; $p = 0.039$); however not for the yaw rate nulling model (panel (e); mean T_R 0.29 s and 0.34 s for experienced and novice drivers; $t(22) = 1.82$; $p = 0.083$). The correlation between T_R and K in panel (e) is statistically significant ($r = -0.58$; $p = 0.003$).

¹A $p < 0.05$ criterion for statistical significance is adopted here.

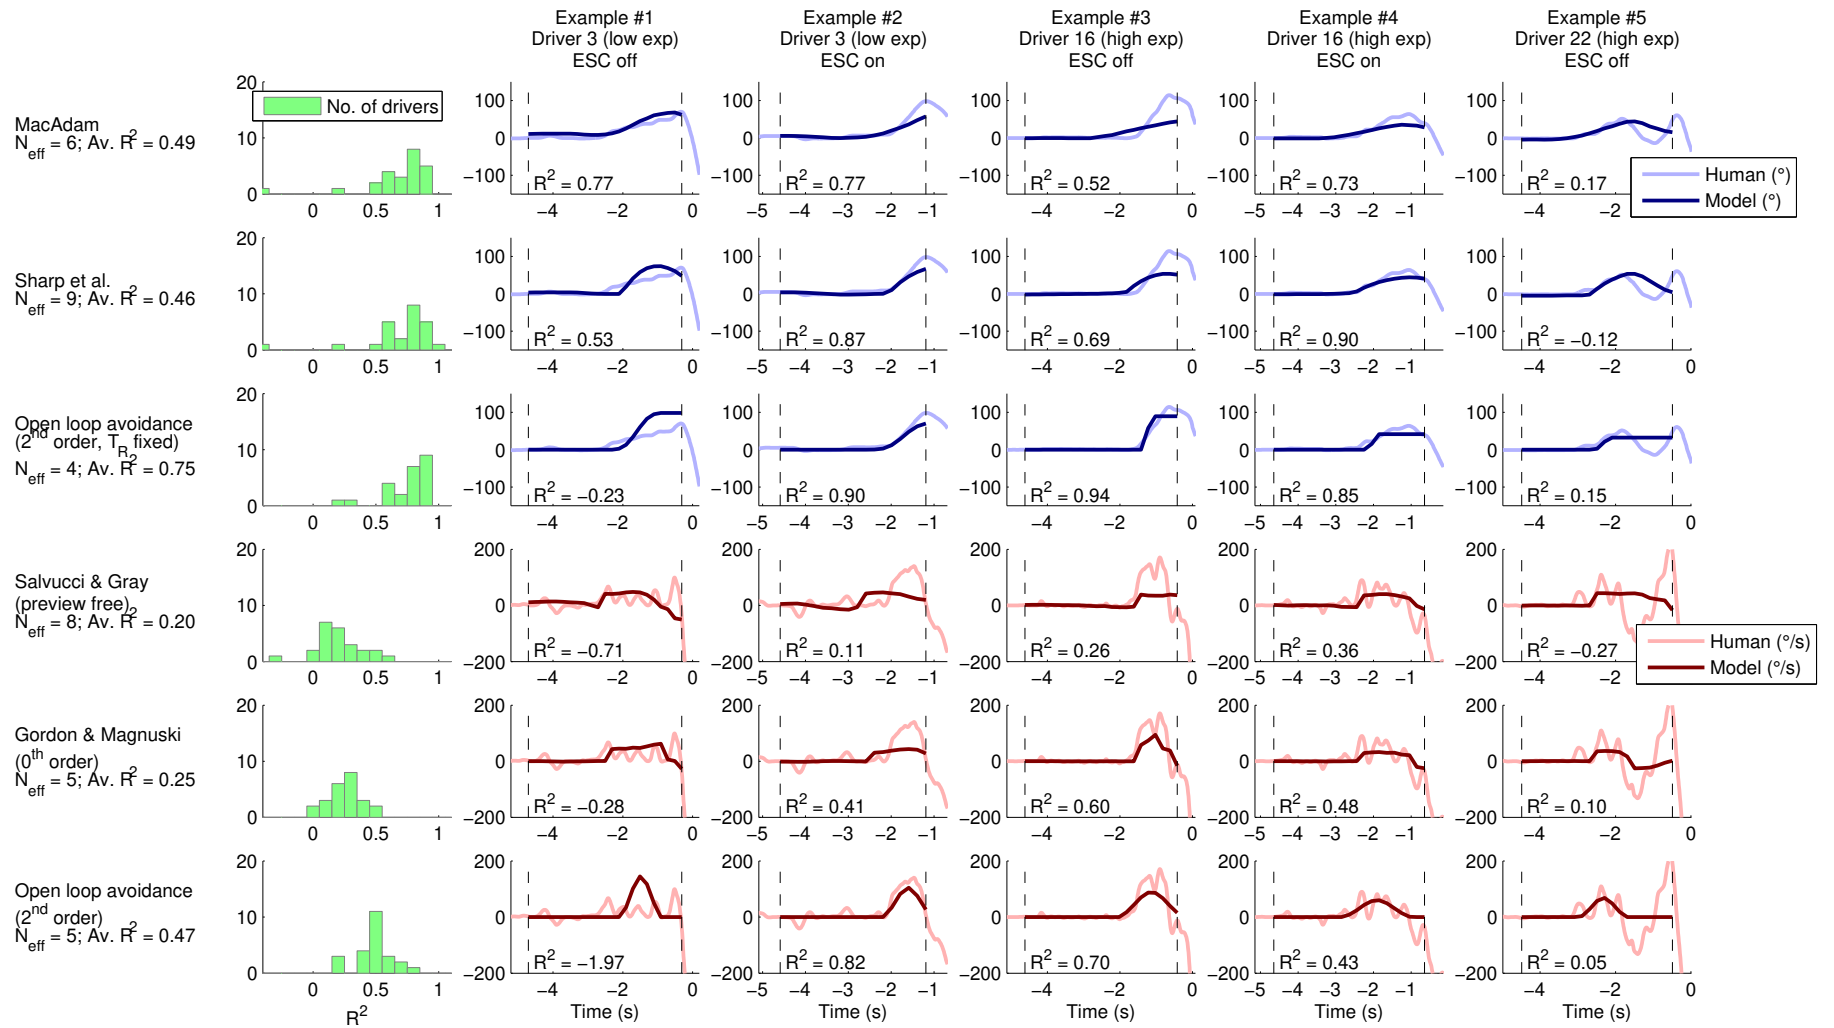


Figure 6. Collision avoidance steering. Distributions of average model performance (leftmost panels) and comparison of human and model steering in five example recorded scenarios, for six of the tested models. Besides including both novices and experienced drivers, driving with and without ESC, the example recordings have been chosen to illustrate typical model performance across various steering behaviours, and to highlight strengths and weaknesses of the different models. Note that the top three rows show models controlling steering angle (δ), whereas the bottom three show models controlling steering rate ($\dot{\delta}$), thus deviating somewhat from the order of model presentation in Sect. 2.2.

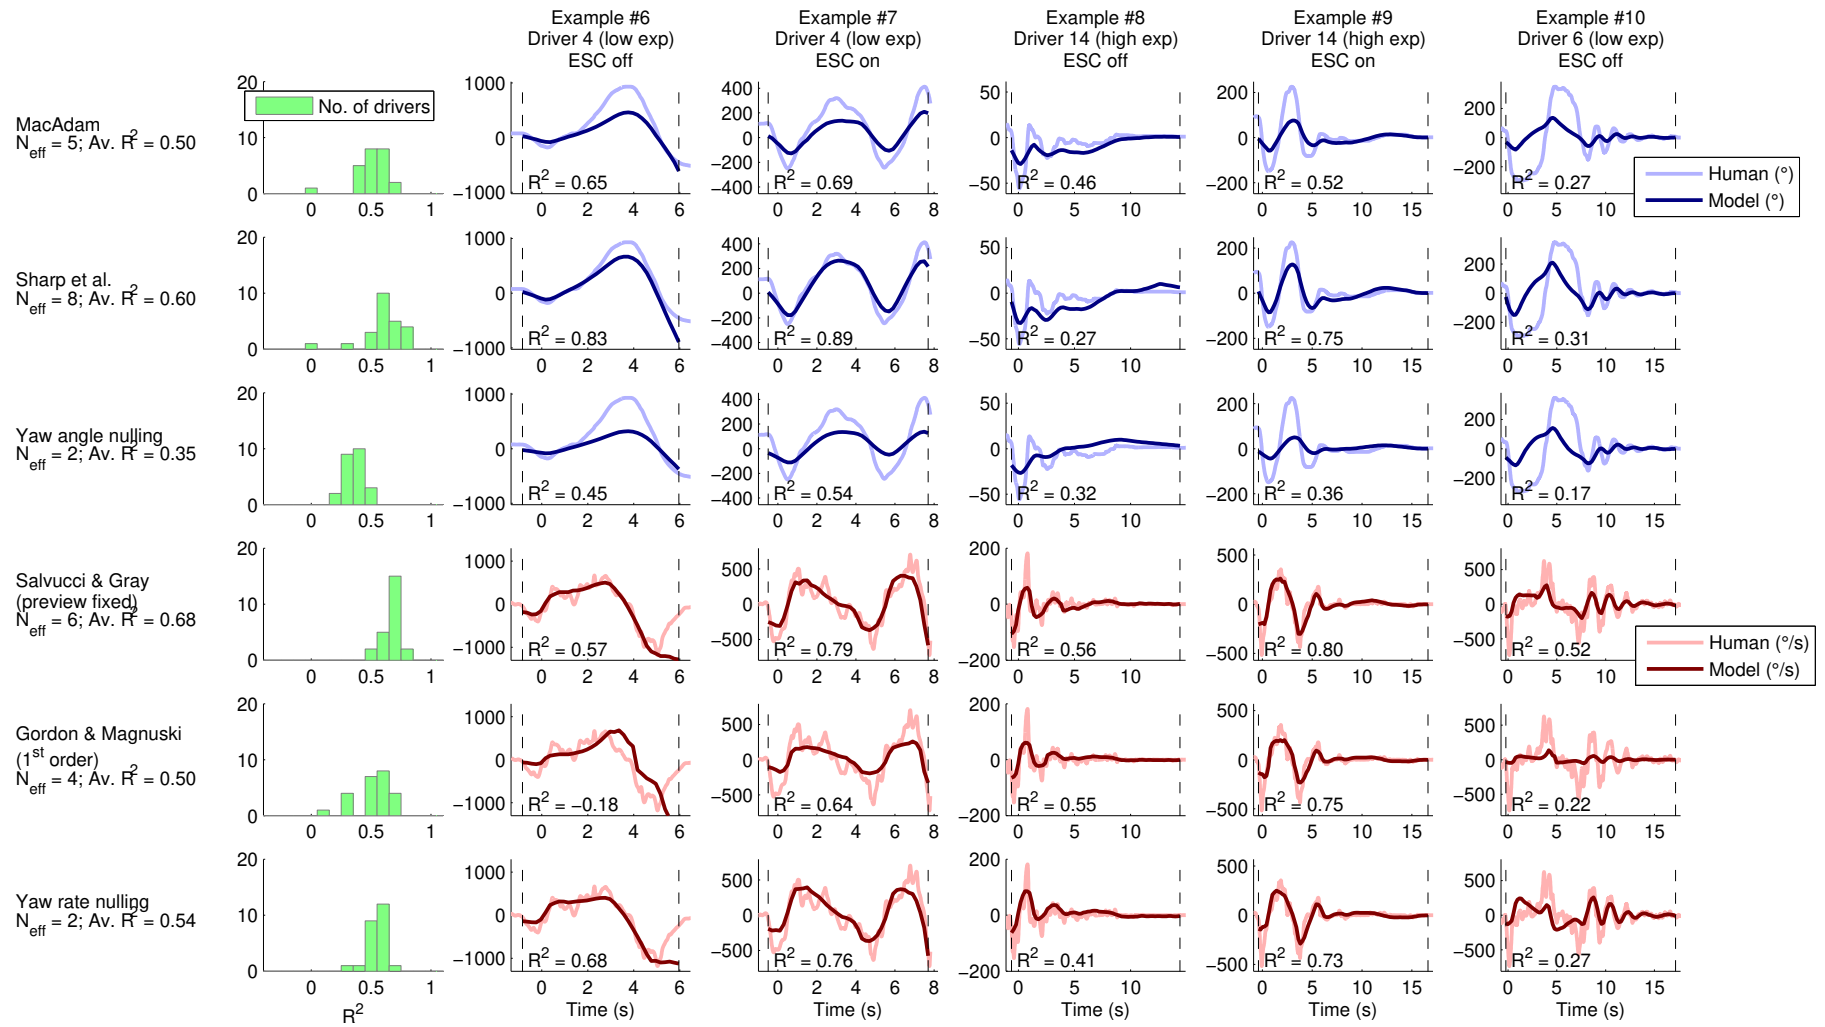


Figure 7. Stabilization steering. Distributions of average model performance (leftmost panels) and comparison of human and model steering in five example recorded scenarios, for six of the tested models. Besides including both novices and experienced drivers, driving with and without ESC, the example recordings have been chosen to illustrate typical model performance across various steering behaviours, and to highlight strengths and weaknesses of the different models. Note that the top three rows show models controlling steering angle (δ), whereas the bottom three show models controlling steering rate ($\dot{\delta}$), thus deviating somewhat from the order of model presentation in Sect. 2.2.

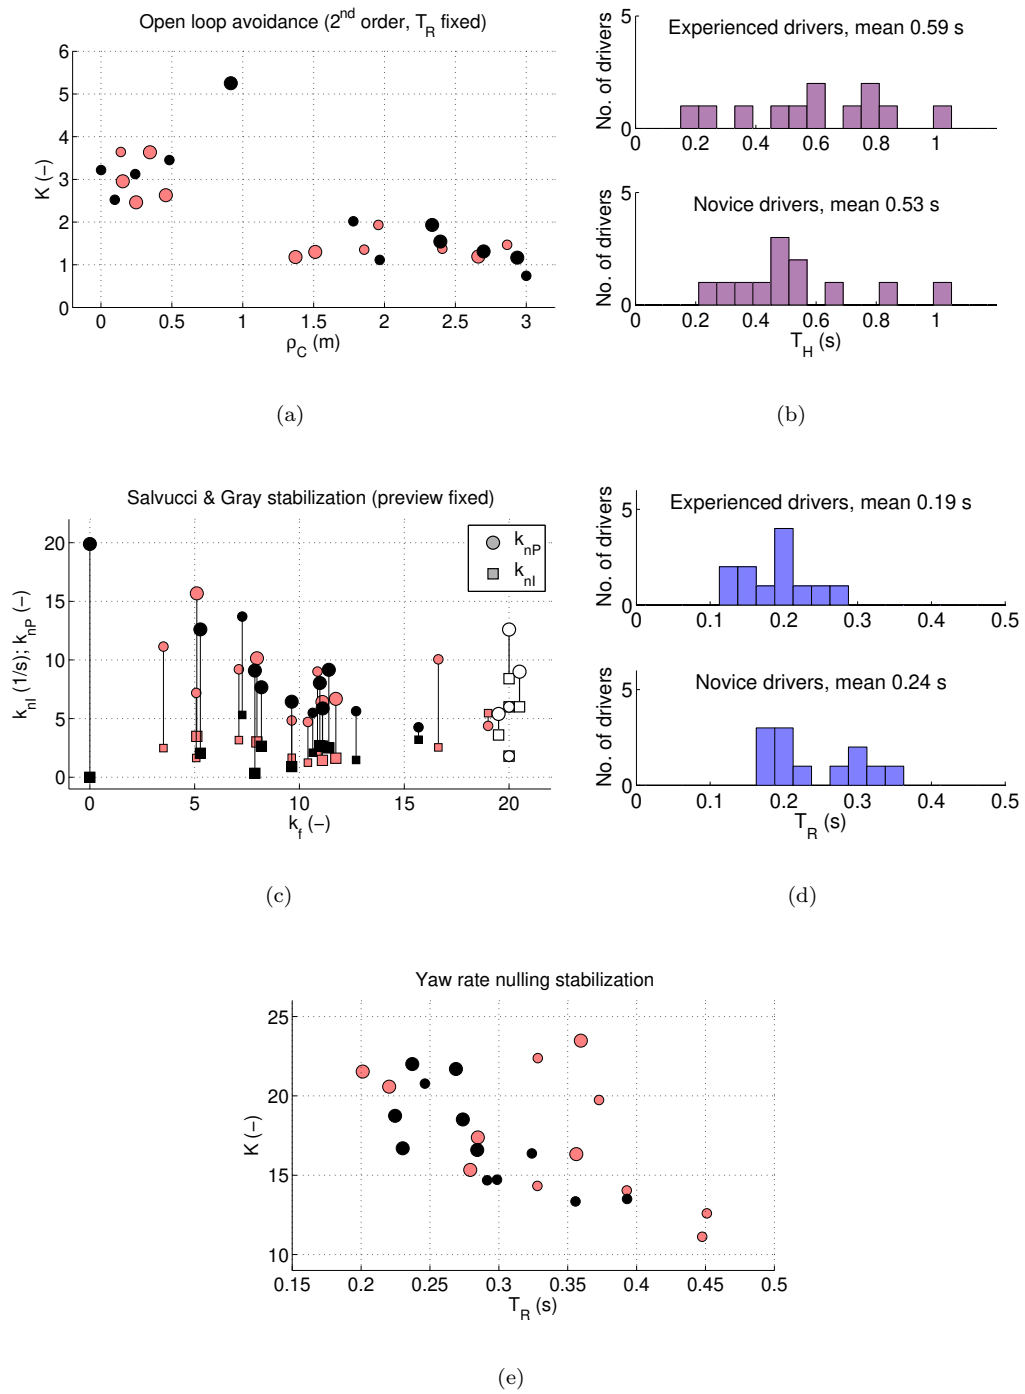


Figure 8. Obtained parameter values for the parameter-reduced variants of the open loop avoidance (panels (a) and (b)) and Salvucci & Gray (panels (c) and (d)) models, as well as for the yaw rate nulling model (panel (e)), for the 24 drivers in the data collection. In panels (a), (c), and (e), black symbols denote experienced drivers, and light red symbols denote novice drivers. Large and small symbols denote drivers for which average model R^2 was above and below the median value among the drivers, respectively. In panel (c), each driver is represented by a pair of one circle and one square, joined by a vertical line. The white symbols show parameterizations suggested for passenger car driving by Salvucci & Gray [37]; all of these are $k_f = 20$, but have been slightly displaced for clarity. In panels (b) and (d), distributions of model parameters T_H (the half-length of the open loop avoidance steering pulse) and T_R (the reaction time of the Salvucci & Gray model) are shown, separately for experienced and novice drivers.

4. Discussion

In general, the very simple open loop avoidance and yaw rate nulling models turned out to work rather well, and the performance of the more advanced models can to some extent be understood as being dependent on an ability to generate the behaviour of the simpler models. These aspects will be discussed below, separately for avoidance and stabilization. Before concluding, some remarks will also be made regarding the challenges involved in comparing and validating driver behaviour models.

4.1. Collision avoidance

4.1.1. Open loop avoidance steering

The open loop avoidance model provided the best fits of the human avoidance steering, both when comparing among models targeting δ and those targeting $\dot{\delta}$. Although there were cases where the human avoidance steering was more gradual (Example #1 in Fig. 6) or oscillatory (Example #5), in a majority of cases most of the total steering angle change was applied in a short period of time (Examples #2–#4), such as suggested by the correlation in Fig. 5(a). This type of open-loop account of collision avoidance steering is not new, and similar models have been used not the least in accident reconstruction work and *what-if* simulations [46–48].

The difference in fit between the constant and variable amplitude variants of the model (especially notable for the δ -controlling model; $R^2 = -1.20$ versus R^2 between 0.54 and 0.76) implies that drivers adapted their avoidance steering to the specific situation. The highest observed R^2 values (0.76 and 0.47, for the δ and $\dot{\delta}$ controlling variants, respectively) were obtained with the assumption that drivers selected their avoidance amplitude based on a 2nd order steering requirement prediction. This suggests that drivers may have been able to take the non-zero deceleration of the lead vehicle into account. On the other hand, the model variants based on looming, which does not include any acceleration information, reached almost as high R^2 values (0.71 and 0.46), so the results are far from conclusive in this respect.

When T_R was a free parameter in the optimizations, it varied throughout the entire permitted optimization interval of $[0, 2]$ seconds, something which could be interpreted as this parameter not being highly important for achieving a good fit, and this is confirmed by the very minor decrease in validation R^2 (from 0.76 to 0.75) when fixing T_R at 0.2 s. It is this parameter-reduced variant of the model which is the basis of Figs. 8(a) and (b). In (a), the correlation between ρ_C and K is clearly due to two clusters of parameterizations. These are interpretable, respectively, as (1) steering roughly as deemed necessary given the linear, low-friction vehicle model (K close to 1) to achieve what seems like an unrealistically large safety margin of about 1.5 – 3 m, and (2) steering aiming for a small safety margin of about 0 – 0.5 m, but applying a steering about three times larger than what the linear vehicle model predicts would be needed for this purpose. A possible interpretation of these fits is that the drivers adapted to the low friction circumstances, responding to vehicle understeering by applying larger steering angles than they would normally [23, 49]. Indeed, a separate, cursory analysis of the avoidance steering in the first, unexpected scenario suggests that while steering was predominantly pulse-like already at this point in the experiment (as implied also by Fig. 5(a)), the steering pulse amplitudes were generally smaller than predicted by the models fitted to the repeated-scenario data. To further clarify exactly how drivers select

their avoidance steering amplitude, more data would be needed, from more varied kinematic situations.

With regards to the duration of avoidance steering (parameter T_H), Fig. 8(b) shows considerable variation between individuals, between 0.2 s and 1 s. The steering rate plots in the bottom three rows of Fig. 6 suggest that this is due to variations in the number of smaller steering corrections needed to achieve satisfactory collision avoidance (cf. [50]). It can also be noted that the average of 0.56 s shown in Fig. 8(b) is not far from the 0.42 s predicted by the slope of the correlation in Fig. 5. That these values are not exactly identical despite being estimated from the same data set is not surprising, if one considers the major differences between the two methods of estimation.

4.1.2. The other models of avoidance steering

In the cases where most of the steering wheel change occurred in a brief period of time (especially clear in Examples #2 and #3 in Fig. 6), the closed loop δ -controlling models by Salvucci & Gray and Gordon & Magnuski were less able than the δ -controlling open loop model at reproducing the resulting overall pulse of steering change. The closed loop models controlling δ (MacAdam and Sharp *et al.*) did produce the corresponding step-like δ outputs, but reached lower average validation R^2 than the δ -controlling open loop model, despite having a higher number of free parameters.

Besides lower R^2 values, another possible objection to the MacAdam and Sharp *et al.* models is related to their use of the desired path construct, which in the context of collision avoidance could be seen as problematic in at least two ways: (1) With a moving lead vehicle, the desired path will, during a first period of time, typically pass *through* the lead vehicle, which makes this construct less attractive than it may seem in scenarios where a path can be charted between stationary obstacles or lane boundaries. (2) There is a parameter redundancy by which an entire single lane change path (such as in this collision avoidance scenario) can be shifted longitudinally without affecting the steering behaviour, as long as one or both of the preview and reaction time parameters are appropriately modified at the same time. Besides these specific issues, it can also be noted that recent neurobiological models of basic sensorimotor control seem to be moving away from desired trajectory constructs, instead placing emphasis on *goal states* [51, 52], arguably more similar to the target lateral position of the Salvucci & Gray model or the Gordon & Magnuski model's goal of avoiding obstacles and lane boundaries.

4.2. Stabilization

4.2.1. Yaw rate nulling stabilization steering

When it comes to stabilization steering, it has been previously shown that models fitted to the repeated scenario data could successfully predict also unexpected scenario behaviour [32]. Here, the most important new result is the good fits obtained for the yaw rate nulling model. Given that the highest average validation R^2 across all stabilization models was 0.68, for the Salvucci & Gray model with six or eight free parameters (preview distances fixed or free), the R^2 of 0.54 yaw rate nulling model, with only two free parameters, seems very good. Qualitatively, the fits (such as shown in Fig. 7) are also rather convincing. The most natural interpretation of these observations seems to be that in the studied scenario, stabilization steering was indeed driven to a large extent by a control law similar to what the yaw rate nulling model suggests.

Three main types of cases were identified where the yaw rate nulling model did

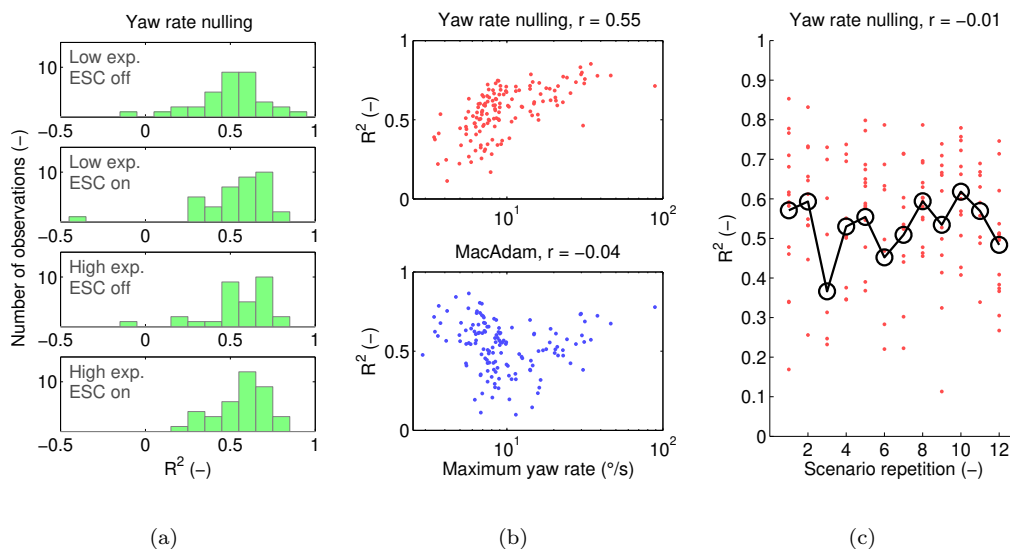


Figure 9. A more detailed view of model fits for the yaw rate nulling and MacAdam models, as a function of (a) experimental conditions, (b) maximum yaw rate attained during each scenario, and (c) scenario repetition. In panels (b) and (c), each small dot corresponds to one recorded scenario from the model-fitting validation set (three points with $R^2 < 0$ not shown for the yaw rate nulling model, two for the MacAdam model), and the r values are the corresponding Pearson correlation coefficients. In (a), the r values were calculated with the logarithm of the maximum yaw rates (such as shown here). Without the logarithms, $r = 0.40$ (top) and $r = 0.05$ (bottom). In (b), the rings are averages per repetition.

not work well: (1) Cases where the driver seemingly gave up steering in the face of imminent control loss (Example #6 in Fig. 7), possibly accounting to some extent for the left tails of the non-ESC distributions in Fig. 9(a), since control losses were more common without ESC [33]. (2) Cases with less vehicle instability and less critical steering (Example #8 in Fig. 7). Fig. 9(b) illustrates this phenomenon in more detail, by showing increasing model fits for increasing maximum vehicle yaw rates. (3) One or two novice drivers (including driver 6, see Example #10 of Fig. 7), who seem to have been using steering strategies of a qualitatively different kind, possibly accounting for the low-experience distributions in Fig. 9(a) being marginally farther to the left (combined average $R^2 = 0.52$) than the high-experience distributions (combined average $R^2 = 0.56$).

The pattern of lower steering gains K in drivers with longer reaction times T_R , shown in Fig. 8(e), can be interpreted as an adaptation of steering aggressivity to one's own response speed, to ensure vehicle stability. Such adaptation could have occurred as a learning effect during the experiment, but the lack of any clear effect of scenario repetition on model fit (Fig. 9(c)) rather suggests that drivers came to the experiment with this adaptation already in place.

4.2.2. The other models of stabilization steering

However, the yaw rate nulling model can hardly provide a full account of steering in the studied scenario; it can stabilize a vehicle directionally, but it has no means to make it stay on a road or close to some path. In contrast, all of the other tested models have such means, and all of them can also be made, more or less naturally, to exhibit some degree of yaw rate nulling.

The yaw angle nulling model will, by definition, have a steering rate of the yaw rate nulling form. Nevertheless, it provides rather poor fits of the human steering angle data (average validation $R^2 = 0.35$). This could possibly be due to the model's lack of a desired path or similar construct, making it unable to

exhibit the rightward lane change during stabilization¹. To keep the corresponding low-frequency component error down (e.g. in the second halves of Examples #8 and #9 in Fig. 7), the optimization may have favoured lower steering gains, in turn making the model unable to generate high-frequency steering of sufficient magnitudes during vehicle instability.

If so, the better fits of the Sharp *et al.* model (average validation $R^2 = 0.60$) could be due to its yaw angle nulling (the term in e_ψ) being relative to a desired path, but the lateral position error terms in the model may of course also have contributed. A main source of reduced R^2 values for this model seems to have been cases where the driver deviated from his or her own typical path in the scenario, such that the model's fitted desired path was not appropriate (Example #8).

The MacAdam model also has a desired path, but no direct means of applying yaw angle or yaw rate nulling. The fits obtained here had long preview times T_P (average 3.6 s, compared to 2.6 s for the Sharp *et al.* model). This makes the optimal control prioritize following the general direction of the road over correcting for local lateral position errors, in essence reducing it to the yaw angle nulling model, a similarity which is clear from Fig. 7. The resulting validation R^2 average of 0.50 was slightly lower than for the yaw rate nulling model, despite the larger number of free parameters, and there was no increase in model fit with increasing yaw instability (Fig. 9(b)).

The Gordon & Magnuski model, by Eq. (5), applies yaw rate nulling as long as $|\dot{\psi}_{\text{req}}| \ll |\dot{\psi}|$. However, when close to a lane exceedence (Example #6 in Fig. 7) the model prioritizes lane keeping higher than the human drivers did. Overall, the model is also less aggressive in its yaw rate nulling behaviour than the humans (see the other examples in Fig. 7); possible reasons for this include the steering gains being kept down (the τ_s being kept high) to minimize error when yaw rates are low, and the model's satisficing approach of aiming for non-zero yaw rate remainders where the humans seemingly did not.

As shown mathematically in [32] and illustrated in Fig. 10, on a straight road far point rotation approximates negative yaw rate, such that the far point control of the Salvucci & Gray model can be understood as yaw rate nulling¹. This insight helps explain the success of the Salvucci & Gray model in fitting the stabilization steering data, and also provides a candidate for a perceptual cue supporting yaw rate nulling behavior. However, the fact that the far point was parameter-fitted, here, to $D_f = 123$ m ahead of the truck, whereas the 3° down from the horizon suggested by previous authors [37, 40] correspond to $D_f \approx 50$ m for the truck in the experiment, could be taken to suggest that also other cues, such as vestibular cues [23] or large-field visual motion [55] may have been at play.

With regards to the other parameters of the Salvucci & Gray model, it is interesting to note the statistically significant faster response times for experienced drivers (Fig. 8(d)), in line with what has been suggested by several other authors; see e.g. [23] and [20, pp. 1132–1133]. The correlation between k_f and k_{nP} (Fig. 8(c)) could be understood as a parameter redundancy; one which is not surprising given the strong correlation between near and far point rates visible in Fig. 10. It is clear that with $D_n = 16$ m, also the near point angle rate was a very close approximation of negative yaw rate, especially at low lateral speeds relative to the road. The exact values obtained here for k_f and k_{nP} should therefore not be attributed too much importance; they could simply be a more or less arbitrary division of the yaw rate nulling model's single gain parameter K .

¹For studies of a yaw angle nulling model with a desired path, see [53] or [54].

¹On a circular road, far point rotation nulling corresponds to nulling of yaw rate error.

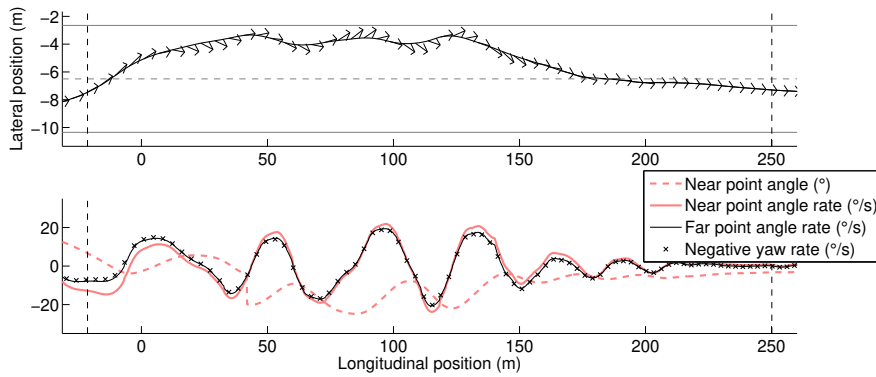


Figure 10. An illustration of the input quantities used by the Salvucci & Gray [37] model, in one example recording. The top panel shows the vehicle trajectory, including arrows showing momentary heading of the truck's front. The bottom panel shows the model input quantities, for $D_n = 16$ m and $D_f = 123$ m, as well as negative vehicle yaw rate. The discontinuity in the near point angle plot corresponds to the truck reaching longitudinal position X_3 , where the model switches its target lane to the right. In both panels, longitudinal position zero corresponds to the point where the truck's front reached the rear of the lead vehicle, and the vertical dashed lines show the beginning and end of the stabilization phase as defined in Sec. 2.3.

4.3. Comparing models of driver control behaviour

One limitation in the comparisons presented here arises from some of the models predicting steering wheel angle δ , and others its time derivative $\dot{\delta}$. The differentiation from δ to $\dot{\delta}$ attenuates steering variations at low frequencies and amplifies those at high frequencies, which means that model-fitting to these two signals will put emphasis on different aspects of steering. For example, the δ and $\dot{\delta}$ -controlling variants of the open loop avoidance model are logically equivalent, but the limitations of assuming a single burst of steering are more obvious in the $\dot{\delta}$ signal than in the δ signal, and this results in lower R^2 values for the $\dot{\delta}$ model variants. Indeed, none of the tested $\dot{\delta}$ -controlling models include any input signals or mechanisms which could have allowed them to fully reproduce the type of high-frequency variations in $\dot{\delta}$ visible in Figs. 6 and 7.

Another limitation, here, is the informal treatment of model parameter count. In theory, additional model parameters cannot reduce model fit, only increase it, and will at the same time increase the risk of obtaining parameter values which overfit to regularities that are unique to the specific data set at hand (e.g. due to only considering one single driving scenario, such as here), thus potentially reducing generality of the parameter-fitted model. There are statistical methods for properly managing this trade-off between model complexity and model fit [56], but these are devised for probabilistic models, as opposed to the completely deterministic models considered here.

Because of the limitations outlined above, the results presented here do not provide grounds for a conclusive recommendation on what model or models to prefer for e.g. simulated evaluation of ESC. Leaving between-model R^2 comparisons to the side, an advantage of the Sharp *et al.* model is that it performed reasonably well both in avoidance and stabilization, implying that one could use a single model for an entire scenario. On the other hand, it could be argued that the Salvucci and Gray model seems more psychologically plausible, due to its input quantities being readily available to a human driver, and since it does not need to assume an internally planned desired path. Psychological plausibility may not be a major priority in some applied contexts, but could provide the benefit of a model that generalizes better beyond the specific data to which it has been parameter-fitted.

It should also be acknowledged, however, that while parameter-fitting of models

Table 2. A summary of advantages and disadvantages of the various tested models, in the studied low-friction collision avoidance scenario.

Model	Advantages in the studied scenario	Disadvantages in the studied scenario
MacAdam	Vehicle-independent parameters. Reasonable fit of avoidance steering.	Relies on desired path construct to achieve fit of avoidance steering. Rather poor fit of stabilization steering.
Sharp et al.	Reasonable fits of both avoidance and stabilization steering.	Relies on desired path construct. Many parameters.
Salvucci & Gray	Best fit of stabilization steering. Psychologically plausible input quantities.	Poor fit of avoidance steering. Rather many parameters.
Gordon & Magnuski	Psychological plausibility of satisficing approach.	Poor fit of avoidance steering and rather poor fit of stabilization steering.
Open loop avoidance	Best fit of avoidance steering.	Not applicable to stabilization steering.
Yaw angle nulling	Few parameters.	Not applicable to avoidance steering. Poor fit of stabilization steering. Not useful as a closed-loop model.
Yaw rate nulling	Reasonable fit of stabilization steering with few parameters, thus potentially indicative of a relevant behavioural phenomenon.	Not applicable to avoidance steering. Not useful as a closed-loop model.

such as performed here may be useful for understanding differences between models, and for pruning out models which do not work at all, it is not necessarily a suitable method for elucidating psychological mechanisms [57]. Here, the good fits of the yaw rate nulling model seem rather compelling, since this model has so few free parameters, but due to the parameter count effects discussed above, the higher R^2 values for the six-parameter Salvucci & Gray model should not be taken as *proof* of that model's underlying assumptions, e.g. that drivers are using near and far points to guide their steering. To study underlying mechanisms, a better approach is to instead identify situations where competing models diverge in their predictions about human behaviour, and then test these predictions in experiment [57].

In future model comparisons, to avoid the δ - $\dot{\delta}$ type of difficulty, one could consider fitting all models to the same control signal (e.g. $\dot{\delta}$). A study of closed-loop behaviour of the parameter-fitted models is also a natural next step, but has been beyond the scope here.

5. Conclusion

The work presented here has clarified some similarities and differences between a number of existing and novel models of driver steering. The strengths and weaknesses of these models in the specific studied scenario are summarized in Table 2. While it has been shown that several of the tested models were reasonably capable of reproducing the observed human steering behaviour, it also remains clear that, even within a well-defined and constrained context, it is non-trivial to decide which exact models to prefer over others. Furthermore, the poor fits reported here for some models do not imply that these models cannot work well in other contexts or scenarios. Especially regarding the Salvucci & Gray and Gordon & Magnuski models, it should be acknowledged that they were originally formulated for routine lane keeping, rather than near-limit manoeuvring.

Overall, model fits were not much affected by whether drivers were novices or experienced, or whether they were driving with ESC on or off. The drivers included here were all truck drivers, driving a simulated truck, but the simplicity of the main observed behavioural phenomena (open loop avoidance and yaw rate nulling

stabilization) makes it reasonable to hypothesize that these phenomena could occur also in passenger car driving.

Regarding the general approach to modelling human control behaviour, the various models tested here are based on rather different underlying assumptions: From the MacAdam model that emphasizes an internalized model of vehicle dynamics and a desired path, to the Salvucci & Gray model that instead emphasizes visual cues allowing the driver to aim for a target lateral position. The results and analyses presented here show that these types of accounts can predict equivalent behaviour in some cases, but could diverge in others:

Collision avoidance steering was best described by the open-loop model, and the best-fitting version (average validation $R^2 = 0.76$) could be interpreted as the drivers acquiring an updated internal model of the vehicle's behaviour on the low-friction road. However, the version of the same model based instead on optical expansion of the lead vehicle on the driver's retina performed almost as well (average validation $R^2 = 0.71$), with one free parameter less.

Stabilization steering was explained to a large extent (average validation $R^2 = 0.54$), and especially well in recordings with more pronounced yaw instability, by the two-parameter yaw rate nulling model, according to which drivers apply a steering rate proportional to the negative of the vehicle's yaw rate. It is interesting to note that the Salvucci & Gray model (and to some extent also the Sharp *et al.* and Gordon & Magnuski models) clearly does predict a causation from instability to yaw rate nulling, due to sight point rotation nulling, and the Salvucci & Gray model also provided the highest average validation R^2 of 0.68 for the stabilization steering data.

The MacAdam model on the other hand, despite having more free parameters than the yaw rate nulling model, provided slightly worse fits of the stabilization data (average validation $R^2 = 0.50$, without a trend of better fits for more pronounced yaw instability). While, again, it is possible that better fits could be obtained by assuming that drivers acquired a more advanced, non-linear internal vehicle model [23, 49], compensating for tyre saturation with increased steering, it is presently not clear whether additional layers of assumed driver insight into vehicle dynamics would in the end really result in such a simply described, and seemingly non-optimal, behaviour as yaw rate nulling.

The fact that yaw rate nulling behavior during instability is correctly predicted by models originally devised for non-critical driving suggests the interesting possibility that drivers may be applying the same sensorimotor control heuristics (e.g. sight point rotation nulling) in both routine and critical situations (cf. [58]). By such an account, seemingly optimal, vehicle-dynamics-adapted behaviour from drivers in routine driving situations can be understood as these sensorimotor heuristics, although far from optimal in general, being precisely tuned for performance and efficiency after extended practice in a constrained operating regime. This would imply that modellers can afford themselves the practical advantages of optimal control theory and internal vehicle dynamics representations, when simulating driving situations of which the modelled driver has much experience (e.g. normal driving for normal drivers, race car driving for race car drivers). However, when modelling less frequent situations, such as traffic near-crashes, one may be better off with a model that is based on the underlying sensorimotor heuristics.

Acknowledgments

This work was supported by VINNOVA (the Swedish Governmental Agency for Innovation Systems; grant number 2009-02766). The data collection was also sup-

ported by the competence centre ViP (Virtual Prototyping and Assessment by Simulation, www.vipsimulation.se, co-financed by the VINNOVA grant 2007-03083 and the competence centre partners).

References

- [1] J.R. Treat, N.S. Tumbas, S.T. McDonald, D. Shinar, R.D. Hume, R.E. Mayer, N.J. Stansifer, and N.J. Castellan, *Tri-Level Study of the Causes of Traffic Accidents*, Final Report DOT HS 805 099, US Department of Transportation, 1979.
- [2] J.D. Lee, *Fifty years of driving safety research*, Human Factors: The Journal of the Human Factors and Ergonomics Society 50 (2008), pp. 521–528.
- [3] T. Dingus et al., *The 100-car naturalistic driving study, Phase II - Results of the 100-car field experiment*, Interim Report DOT HS 810 593, Virginia Tech Transportation Institute, 2006.
- [4] M. Räsänen and H. Summala, *Attention and expectation problems in bicycle-car collisions: An in-depth study*, Accident Analysis & Prevention 30 (1998), pp. 657–666.
- [5] M. Green, *"How long does it take to stop?" Methodological analysis of driver perception-brake times*, Transportation Human Factors 2 (2000), pp. 195–216.
- [6] G. Malaterre, F. Ferrandez, D. Fleury, and D. Lechner, *Decision making in emergency situations*, Ergonomics 31 (1988), pp. 643–655.
- [7] L.D. Adams, *Review of the literature on obstacle avoidance maneuvers: braking versus steering*, UMTRI-94-19, The University of Michigan Transportation Research Institute, 1994.
- [8] D.V. McGehee, E.N. Mazzae, G.H.S. Baldwin, P. Grant, C.J. Simmons, J. Hankey, and G. Forkenbrock, *Examination of Drivers' Collision Avoidance Behavior Using Conventional and Antilock Brake Systems on the Iowa Driving Simulator*, Final report, University of Iowa, 1999.
- [9] S.E. Lee, E. Llaneras, S. Klauer, and J. Sudweeks, *Analyses of rear-end crashes and near-crashes in the 100-car naturalistic driving study to support rear-signaling countermeasure development*, DOT HS 810 846, Virginia Tech Transportation Institute, 2007.
- [10] W. Levison et al., *Development of a Driver Vehicle Module for the Interactive Highway Safety Design Model*, FHWA-HRT-08-019, Federal Highway Administration, 2007.
- [11] M. Perel, *The development of a computer simulation model of driver performance to predict accident probability*, in *Proceedings of the Human Factors Society 26th Annual Meeting*, 1982, pp. 239–243.
- [12] T. Brown, J. Lee, and D. McGehee, *Human performance models and rear-end collision avoidance algorithms*, Human Factors: The Journal of the Human Factors and Ergonomics Society 43 (2001), pp. 462–482.
- [13] T. Gordon et al., *Advanced Crash Avoidance Technologies (ACAT) Program – Final Report of the Volvo-Ford-UMTRI Project: Safety Impact Methodology for Lane Departure Warning – Method Development and Estimation of Benefits*, Final Report DOT HS 811 405, U.S. Department of Transportation, 2010.
- [14] K.D. Kusano and H.C. Gabler, *Safety Benefits of Forward Collision Warning, Brake Assist, and Autonomous Braking Systems in Rear-End Collisions*, IEEE Transactions on Intelligent Transportation Systems 13 (2012), pp. 1546–1555.
- [15] M. Brackstone and M. McDonald, *Car-following: a historical review*, Transportation Research Part F: Traffic Psychology and Behaviour 2 (1999), pp. 181–196.
- [16] X. Wang, X. Yang, G. Shan, and F. Wang, *Review of the simulation model of driving behavior*, in *Proceedings of the International Conference on Machine Learning and Cybernetics*, 2006, pp. 911–918.
- [17] T. Jürgensohn, *Control theory models of the driver*, in *Modelling Driver Behaviour in Automotive Environments*, in *Modelling Driver Behaviour in Automotive Environments*, ed. P. CacciabueP. Cacciabue ed., Springer, 2007, pp. 277–292.
- [18] M. Plöchl and J. Edelman, *Driver models in automobile dynamics application*, Vehicle System Dynamics 45 (2007), pp. 699–741.
- [19] D.H. Weir and K.C. Chao, *Review of Control Theory Models for Directional and Speed Control*, in *Modelling Driver Behaviour in Automotive Environments: Critical Issues in Driver Interactions with Intelligent Transport Systems*, in *Modelling Driver Behaviour in Automotive Environments: Critical Issues in Driver Interactions with Intelligent Transport Systems*, ed. P. CacciabueP. Cacciabue ed., Springer, 2007.
- [20] G. Markkula, O. Benderius, K. Wolff, and M. Wahde, *A review of near-collision driver behavior models*, Human Factors: The Journal of the Human Factors and Ergonomics Society 54 (2012), pp. 1117–1143.
- [21] K. Lee and H. Peng, *Identification and verification of a longitudinal human driving model for collision warning and avoidance systems*, International Journal of Vehicle Autonomous Systems 2 (2004), pp. 3–17.
- [22] A. Ungoren and H. Peng, *An adaptive lateral preview driver model*, Vehicle System Dynamics 43 (2005), pp. 245–259.
- [23] C. MacAdam, *Understanding and modeling the human driver*, Vehicle System Dynamics 40 (2003), pp. 101–134.
- [24] ISO, *Passenger cars – Test track for a severe lane-change manoeuvre – Part 1: Double lane-change*, .
- [25] J. Breuer, *Analysis of driver-vehicle-interactions in an evasive manoeuvre - Results of 'moose test' studies*, in *Proceedings of the 15th ESV Conference*, 1998, pp. 620–627.
- [26] M. Dilich, D. Kopernik, and J. Goebelbecker, *Evaluating driver response to a sudden emergency: Issues of expectancy, emotional arousal and uncertainty*, 2002-01-0089, SAE International, 2002.

- [27] E. Hollnagel and D.D. Woods *Joint cognitive systems: Foundations of cognitive systems engineering*, CRC Press / Taylor & Francis, 2005.
- [28] A.T. van Zanten, *Bosch ESP Systems: 5 Years of Experience*, Technical Paper 2000-01-1633, SAE International, 2000.
- [29] A. Eichberger, E. Tomasch, R. Rohm, H. Steffan, and W. Hirschberg, *Detailed Analysis of the Benefit of Different Traffic Safety Systems in Fatal Accidents*, in *Proceedings of the 19th EVU Congress, Prague*, 2010.
- [30] S. Moon, W. Cho, and K. Yi, *Intelligent vehicle safety control strategy in various driving situations*, *Vehicle System Dynamics* 48 (2010), pp. 537–554.
- [31] H. Xiao, W. Chen, H. Zhou, and J.W. Zu, *Integrated control of active suspension system and electronic stability programme using hierarchical control strategy: theory and experiment*, *Vehicle System Dynamics* 49 (2011), pp. 381–397.
- [32] G. Markkula, *Evaluating vehicle stability support systems by measuring, analyzing, and modeling driver behavior*; Licentiate thesis, Chalmers University of Technology, 2013.
- [33] G. Markkula, O. Benderius, K. Wolff, and M. Wahde, *Effects of experience and electronic stability control on low friction collision avoidance in a truck driving simulator*, *Accident Analysis & Prevention* 50 (2013), pp. 1266–1277.
- [34] O. Benderius, G. Markkula, K. Wolff, and M. Wahde, *Driver behaviour in unexpected critical events and in repeated exposures - a comparison*, *European Transport Research Review* (2013).
- [35] C. MacAdam, *Application of an optimal preview control for simulation of closed-loop automobile driving*, *IEEE Transactions on Systems, Man, and Cybernetics* 11 (1981), pp. 393–399.
- [36] R. Sharp, D. Casanova, and P. Symonds, *A mathematical model for driver steering control, with design, tuning and performance results*, *Vehicle System Dynamics* 33 (2000), pp. 289–326.
- [37] D. Salvucci and R. Gray, *A two-point visual control model of steering*, *Perception* 33 (2004), pp. 1233–1248.
- [38] T. Gordon and N. Magnuski, *Modeling normal driving as a collision avoidance process*, in *Proceedings of 8th International Symposium on Advanced Vehicle Control*, 2006.
- [39] R.M. Wilkie and J.P. Wann, *Controlling Steering and Judging Heading: Retinal flow, visual direction and extra-retinal information*, *Journal of Experimental Psychology: Human Perception and Performance* 29 (2003), pp. 363–378.
- [40] M. Land and J. Horwood, *Which parts of the road guide steering?*, *Nature* 377 (1995), pp. 339–340.
- [41] H. Summala, *Towards Understanding Motivational and Emotional Factors in Driver Behaviour: Comfort Through Satisficing*, in *Modelling Driver Behaviour in Automotive Environments*, , in *Modelling Driver Behaviour in Automotive Environments*, ed. P. CacciabueP. Cacciabue ed., Springer, 2007, pp. 189–207.
- [42] S. Chang, *A Flexible Hierarchical Model-Based Control Methodology for Vehicle Active Safety Systems*, University of Michigan, 2007.
- [43] D. Lee, *A theory of visual control of braking based on information about time-to-collision*, *Perception* 5 (1976), pp. 437–459.
- [44] M. Wahde *Biologically Inspired Optimization Methods*, WIT Press, Southampton, UK, 2008.
- [45] A. Field *Discovering statistics using SPSS*, 3rd SAGE Publications, 2009.
- [46] M. Araszewski, A. Toor, R. Overgaard, and R. Johal, *Lane change maneuver modelling for accident reconstruction applications*, 2002-01-0817, SAE International, 2002.
- [47] Y. Sugimoto and C. Sauer, *Effectiveness estimation method for advanced driver assistance system and its application to collision mitigation brake system*, in *Proceedings of the 19th International Technical Conference on the Enhanced Safety of Vehicles, Washington DC, United States*, 2005.
- [48] R.M. van Aken, J.W. Zeller, D.P. Chiang, J. Kelly, J.Y. Silberling, R. Dai, P.C. Broen, A.M. Kirsch, and Y. Sugimoto, *Advanced Crash Avoidance Technologies (ACAT) Program - Final Report of the Honda-DRI Team*, DOT HS 811 454, US Department of Transportation, 2011.
- [49] C. MacAdam, *Development of a driver model for near/at-limit vehicle handling*, Final report UMTRI-2001-43, University of Michigan Transportation Research Institute, 2001.
- [50] O. Benderius and G. Markkula, *Evidence for a fundamental property of steering*, in *Proceedings of the Annual HFES Meeting, Chicago, IL, USA, Oct 27-31, 2014*, forthcoming.
- [51] E. Todorov and M. Jordan, *Optimal feedback control as a theory of motor coordination*, *Nature Neuroscience* 5 (2002), pp. 1226–1235.
- [52] D. Franklin and D. Wolpert, *Computational Mechanisms of Sensorimotor Control*, *Neuron* 72 (2011), pp. 425–442.
- [53] A. Reński, *The Driver Model and Identification of Its Parameters*, Technical Paper 980011, SAE International, 1998.
- [54] M. Lin, A. Popov, and S. McWilliam, *Sensitivity analysis of driver characteristics in driver-vehicle handling studies*, in *Driver Behaviour and Training*, , in *Driver Behaviour and Training*, ed. L. DornL. Dorn ed., Ashgate Publishing Ltd, 2003, pp. 263–276.
- [55] N. Saijo, I. Murakami, S. Nishida, and H. Gomi, *Large-Field Visual Motion Directly Induces an Involuntary Rapid Manual Following Response*, *The Journal of Neuroscience* 25 (2005), pp. 4941–4951.
- [56] H. Motulsky and A. Christopoulos *Fitting models to biological data using linear and nonlinear regression*, Oxford University Press, 2004.
- [57] S. Roberts and H. Pashler, *How Persuasive is a Good Fit? A Comment on Theory Testing*, *Psychological Review* 107 (2000), pp. 358–367.
- [58] G. Markkula, *Modeling driver control behavior in both routine and near-accident driving*, in *Proceedings of the Annual HFES Meeting, Chicago, IL, USA, Oct 27-31, 2014*, forthcoming.

Table A1. A listing of all model parameters included in the model-fitting optimizations. For each parameter, the allowed range in the optimizations is indicated, together with information on what models made use of the parameter, and on whether the model was considered effective in the collision avoidance and stabilization steering phases, respectively (i.e. whether or not it was included in the corresponding N_{eff} count in Table 1).

Parameter	Allowed range	Used by models ^a	Collision avoidance	Stabilization
ΔX_1	$[-50, 50]$ m	M, S	x	
X_2	$[-20, 30]$ m	M, S	x	
X_3	$[20, 70]$ m	M, S, S&G	x	x
X_4	$[100, 250]$ m	M, S, G&M		x
Y	$[-1.65, -6.5]$ m	M, S, S&G	x	x
T_P	$[0.1, 5]$ s	M, S	x	x
T_R	$[0, 2]$ s	M, S, S&G, G&M, OLA (except constant amplitude variant), YAN, YRN	x	x
K_ψ	$[0, 100]$	S	x	x
K_1	$[0, 100]$	S	x	x
K_P	$[0, 10]$	S	x	x
ΔX_I	$[-50, 0]$ m	S&G, G&M	x	
D_n	$[0.1, 200]$ m	S&G	x	x
D_f	$[0.1, 200]$ m	S&G	x	x
k_f	$[0, 100]$	S&G	x	x
k_{nP}	$[0, 50]$	S&G	x	x
k_{nI}	$[0, 10]$	S&G	x	x
ρ_C	$[0, 3]$ m	G&M, OLA (steering requirement variants)	x	x
ρ_L	$[-2.25, 3]$ m	G&M	x	x
τ_s	$[0, 10]$ s	G&M	x	x
K	$[0, 20]$	OLA (constant amplitude variant)	x	N/A
K	$[0, 200]$	OLA (other variants)	x	N/A
T_A	$[-0.5, 0.5]$ s	OLA	x	N/A
T_H	$[0.1, 1]$ s	OLA	x	N/A
K	$[0, 100]$	YAN, YRN	N/A	x

^aM: MacAdam; S; Sharp *et al.*; S&G: Salvucci & Gray; G&M: Gordon & Magnuski; OLA: Open loop avoidance; YAN: Yaw angle nulling; YRN: Yaw rate nulling.

Appendix A. Genetic algorithm implementation

In the GA used for model-fitting (see Sec. 2.4), a candidate model parameterization was represented by a GA individual with a genome of length N , the number of free parameters of the model. Each gene was a floating point number in the interval $[0, 1]$, corresponding to a value within the allowed range for the parameter in question. These ranges are listed in Table A1, also showing which parameters were considered effective in the collision avoidance and stabilization phases, respectively.

The GA was configured, in the terminology and notation of [44, pp. 48–55], as follows: population size 100, tournament size 2, tournament selection parameter $p_{\text{tour}} = 0.9$, crossover probability $p_c = 1$, mutation probability $p_{\text{mut}} = 1/N$. Mutation consisted in either (with probability 0.5) randomly choosing a new value from a uniform distribution in $[0, 1]$, or otherwise applying real-number creep from a normal distribution of standard deviation 0.005, after which the new value was bounded to $[0, 1]$. The best individual in a given GA generation was always retained in the next generation (elitism). Initial exploration indicated that model-fitting R^2 values were not very sensitive to the exact GA configuration, and the specific GA parameter settings adopted here were selected based on a criterion of low variability in R^2 estimates across repeated optimizations.

The GA was terminated at completion of generation number $2G$, where G was the last generation in the optimization with an increase in validation fitness. However, all optimizations were allowed a minimum of 300 generations.