Identifiability and Practical Relevance of Complex Car-Following Models

Gunnar Flötteröd, Peter Wagner, and Yun-Pang Flötteröd

Abstract This article looks at car-following models from a deliberately pragmatic perspective: What information about driver behavior can be extracted from a given data set without more or less speculative assumptions about underlying behavioral laws? The objective of this exercise is not to invalidate existing models but to obtain a better understanding of how much (complex) model structure can be revealed/validated from real data.

1 Introduction

The estimation of parameters of a microscopic traffic flow model appears at first glance to be a technically straightforward and well understood procedure. What is not that well understood is the question what is actually revealed by the calibration. Typically, the calibration exercise results in parameters that minimize some distance between model outputs and reality. Some of these parameters have immediate physical meanings: maximum speed, maximum acceleration, and the like. Other parameters are hard to interpret and hence are difficult to validate in hindsight, not even through simple plausibility checks. While sophisticated car-following model specifications abound, much of their added value lies in theoretically being able to explain certain (rare) phenomena, which, however, are of limited relevance if one is interested in estimating, say, a car-following model component for a complex network simulation from real data.

G. Flötteröd (✉)
KTH Royal Institute of Technology, Stockholm, Sweden
e-mail: gunnar.floetteroed@abe.kth.se

P. Wagner · Y.-P. Flötteröd
DLR German Aerospace Center, Cologne, Germany
e-mail: peter.wagner@dlr.de; yun-pang.floetteroed@dlr.de

V.V. Kozlov et al. (eds.), *Traffic and Granular Flow '11*,
DOI 10.1007/978-3-642-39669-4_5,

This article adopts an (almost naive) engineering perspective on the problem in that it starts the analysis by estimating a set of simple linear models from a given data set. If a linear model already explains the data well, there is little reason to complexify the model further. If, however, a linear model fails to explain certain aspects of the data, it still is possible to analyze the residuals in order to obtain data-driven hints of how to improve the model. Clearly, an argument against this approach is that a good model structure, based on physical and/or behavioral considerations, should also have a superior explanatory power. The counter-argument has in essence already been phrased: A simple model with only few, interpretable parameters may result in a slightly inferior fit, but its estimation may be more robust than for a complex model, and its simplicity and interpretability is likely to be a key feature in its practical application.

The remainder of this article investigates the above claims by constraining itself to utmost simple models and exclusively inferring model structure either from exact physical laws or from the data itself. Section 2 estimates and analyzes a set of linear car following models. Section 3 then discusses the implications of this article's findings, in particular with respect to more complex nonlinear model specifications.

2 Estimation and Analysis of Linear Models

The data used in this contribution has been recorded some years ago on a Japanese test track [2]. Nine drivers were following a lead driver, who implemented a certain speed protocol, changing between regimes of constant and varying speeds. Most acceleration values were moderate and in the range $[-3, 3]\,\text{m/s}^2$. The drivers' kinematic states were tracked by means of a differential RTK-GPS system, which allows for centimeter accuracy in positions and centimeter-per-second precision in speeds.

Due to the real-time kinematics in the GPS recording, speeds are independently measured variables (essentially based on the Doppler shift in the signal frequency between vehicle and satellites) and are not derived via numerical differentiation from the positions. Altogether, data from eight different experiments of 25 min each is available, with different drivers and different vehicle sequences per experiment. Within this contribution, only the first experiment is analyzed. The full data set has been used in another study on car-following models and their calibration [1].

2.1 Considered Model Specifications

A number of increasingly complex linear models is calibrated from the data described above. The purpose of this is to understand how much information is contained in the data that can be captured by linear dynamics.

We start from the optimal velocity model (OVM) in discrete time t, which is written as

$$v(t) - v(t-h) = \alpha \cdot (v^*(t-h) - v(t-h)) + \varepsilon(t) \tag{1}$$

where v is the considered driver's speed, v^* is the instantaneous optimal velocity (yet to be defined), $\alpha \in (0, 1)$ defines the speed of the driver's adaptation towards v^*, h is the time step length, and ε is a temporally uncorrelated error term. (The last assumption is convenient in terms of model estimation but needs to be validated and possibly adjusted in future work [3].) Specifically, the following model instances are considered.

Speed synchronization model. This model assumes $v^*(t) = V(t)$, which means that the driver does nothing but follow the speed V of the leading car. The resulting linear model specification has only one independent parameter:

$$v(t) = a_1 v(t-h) + c_1 V(t-h) + \varepsilon(t) \tag{2}$$

where a_1 corresponds to $1-\alpha$ and c_1 is constrained to be $1-a_1$.

Newell-type model. This model assumes $v^*(t) = \frac{1}{\tau} g(t)$ with $\tau > 0$, which means that the desired speed is proportional to the distance g from the leading car. The resulting linear model specification has two parameters:

$$v(t) = a_1 v(t-h) + b_1 g(t-h) + \varepsilon(t) \tag{3}$$

where a_1 corresponds to $1-\alpha$ and b_1 corresponds to $\alpha \frac{1}{\tau}$.

Combined OVM model. This model assumes $v^*(t) = \gamma \frac{1}{\tau} g(t) + (1-\gamma) V(t-h)$, $0 \leq \gamma \leq 1$, which defines the desired speed as a convex combination of lead car speed and distance. The resulting linear model specification has three parameters:

$$v(t) = a_1 v(t-h) + b_1 g(t-h) + c_1 V(t-h) + \varepsilon(t) \tag{4}$$

where a_1 corresponds to $1-\alpha$, b_1 corresponds to $\alpha \frac{1}{\tau} \gamma$, and $c_1$corresponds to $\alpha(1-\gamma)$.

Generalized OVM model. A natural extension of the combined OVM model is to allow for higher order dynamics in the speed adaptation process. This leads to the following linear model specification:

$$v(t) = \sum_{i=1}^{\cdots} a_i v(t-ih) + \sum_{i=1}^{\cdots} b_i g(t-ih) + \sum_{i=1}^{\cdots} c_i V(t-ih) + \varepsilon(t). \tag{5}$$

Since this model allows for arbitrary linear combinations of past gaps and velocities of the modeled car and the lead car, it implicitly also captures speed differences as explanatory variables. Also, since $g(t) - g(t-h) = h \cdot (V(t-h) - v(t-h))$, a linear dependency between g, v, and V can be expected. Indeed, it turns out

Table 1 Estimated parameters for a single driver and the three simple linear models (2)–(4). The values in brackets indicate the statistical error in the parameter estimate, it relates to the last digit(s) of the parameter value, i.e. 0.9721 ± 0.0005 is written as 0.9721(5)

Model	$v(t-h)$	$g(t-h)$	$V(t-h)$	σ	LL	AIC
Speed -sync.	0.9721(5)	–	0.0279(6)	0.0565	23,761	−47,516
Newell	0.9968(1)	0.0021(1)	–	0.0595	22,917	−45,828
Combined	0.9719(6)	0.00127(9)	0.02620(6)	0.0561	23,865	−47,723

that the parameters $b_{i>1}$, $c_{i>1}$ cannot be identified. What worked, and therefore is exclusively considered in the following, is the simplification

$$v(t) = \sum_{i=1}^{\dots} a_i v(t-ih) + b_1 g(t-h) + c_1 V(t-h) + \varepsilon(t). \tag{6}$$

2.2 *Description and Analysis of Estimation Results*

A first set of model estimation results is shown in Table 1. All results are obtained with R's generalized linear model routine glm(). Also shown in this table are the following measures of fit:

- The average squared residual

$$\sigma^2 = \frac{1}{N} \sum_{i=1}^{N} \left(v_i^{\text{sim}} - v_i^{\text{data}}\right)^2 \tag{7}$$

  where N is the number of data points, v_i^{sim} is the ith simulated velocity, and v_i^{data} is its corresponding real data point;
- The log-likelihood

$$\text{LL} = \sum_{i=1}^{N} \ln\left(g\left(v_i^{\text{sim}} - v_i^{\text{data}}\right)\right) \tag{8}$$

  where in the given setting g is a standard normal distribution;
- The Akaike Information Criterion

$$\text{AIC} = 2k - 2\text{LL} \tag{9}$$

  where k is the number of parameters in the model. The AIC takes into account not only the likelihood but also the number of parameters needed to achieve

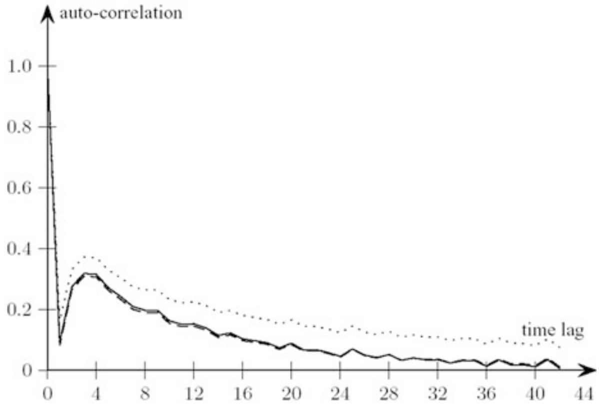

Fig. 1 Auto-correlation function of the residuals. *Solid*: Speed-synchronization, *dotted*: Newell, *dashed*: combined

the corresponding data fit: For two models with the same likelihood values, the model with the smaller AIC can be considered as "better".

Overall, the "combined" model yields both the best fit and the the smallest AIC, indicating its superiority over the two other models. All model parameters are within plausible ranges: the by far largest effect on the instantaneous velocity results from the previous velocity, which reflects the inertia of the driver/vehicle system. Small (but clearly significant) positive coefficients on g and V capture the distance-sensitivity of the driver as well as her adaptation to the preceding vehicle's speed.

Additional insights into the model dynamics can be obtained from Fig. 1, which shows the auto-correlation function of the residuals of the first three models. A first observation is that the Newell model performs substantially worse than the two other models, which is apparently due to the omission of the speed of the leading vehicle.

All models exhibit a relatively low autocorrelation at a lag of one, which then jumps to a larger value again and decays slowly. As a remedy to this "tail" in the auto-correlation, the "generalized" model (6) is implemented with more than one lag in the feedback of v on itself. A maximum lag of seven is chosen (after which the AIC becomes larger again). The resulting coefficients are shown in Table 2. One observes a clearly improved fit and AIC. Also, the obtained parameters are plausible again. The velocity v is fed back with a positive coefficient for small lags, representing inertia. For higher lags, the coefficients become negative, indicating overreaction. This could be (tentatively) interpreted as drivers being unable to react immediately but tending to overreact once their perception is updated.

The error autocorrelation function of the generalized linear model is flat, with only one peak at zero lag, indicating that all *linear* dynamics are captured well by the model.

Table 2 Parameter value(s) for the generalized linear model Eq. (6). All parameters but a_4 have significance levels well below 10^{-6}, the weakest ones being those for $g(t-h)$ and $v(t-5h)$

Model	$v(t-h)$	$g(t-h)$	$V(t-h)$	σ	LL	AIC
Generalized	0.813(8)	0.00037(7)	0.0148(6)			
Higher order	$a_2 = 0.26(1)$	$a_3 = 0.1621$	$a_4 = 0$ (not sign.)	0.0486	26,224	−52,429
Terms	$a_5 = -0.640$	$a_6 = -0.0964$	$a_7 = -0.0871$			

Table 3 Estimation results for the "combined" model, for all drivers

Driver	$v(t-h)$	$g(t-1)$	$V(t-h)$
1	0.9719	0.00127	0.02620
2	0.9748	0.00183	0.02230
3	0.9770	0.00254	0.01897
4	0.9740	0.00102	0.02375
5	0.9757	0.00307	0.01993
6	0.9822	0.00071	0.01648
7	0.9777	0.00159	0.01934
8	0.9679	0.00370	0.02586
9	0.9737	0.00081	0.02471

2.3 Investigation of Driver Heterogeneity and Time-Dependent Parameters

Table 3 shows the estimation results for all drivers individually, using the "combined" model (4). The is very little variation in the obtained parameter values between drivers: the results are extremely stable.

This remains the case when looking at the time-resolved parameter estimates for a single driver. For this, a window of a certain size w is moved along the data, and within each window the "combined" linear model is estimated. As the time window is moved along the data, a time-series of each of the parameters of the linear model is generated. For a window size of $w = 20$ s, the parameters of a single driver again turn out to be fairly stable, with the exception of the beginning of the time series, where the vehicles are standing still and not yet following each other; see Fig. 2 for an example.

2.4 Simulation Results

Now, the space-time plots of the real data are compared with those obtained by simulation of the linear models. To begin with, Fig. 3 shows the space-time plots of the original data for the leading and all nine following vehicles. This plot alone reveals little but the fact that the leading vehicle indeed varies its speed.

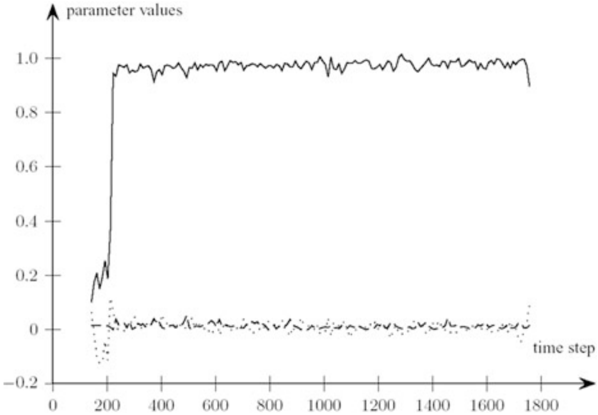

Fig. 2 Time dependence of the parameters of the generalized linear model for driver one. *Solid*: a_1, *dashed*: b_1, *dotted*: c_1

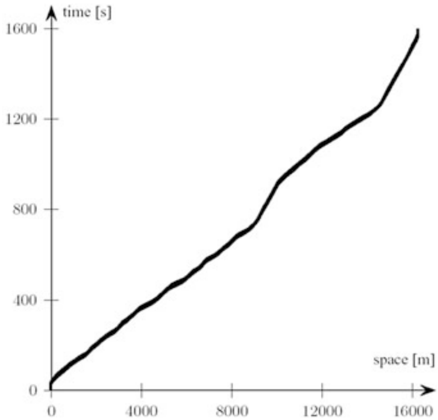

Fig. 3 Space-time plot of original data

Figure 4 (left) shows a "differential" perspective on the original data, where only the following vehicles are plotted relatively to the leading vehicle's coordinates. One clearly recognizes a widening and shortening of the platoon in reaction to the lead vehicle's driving pattern. Figure 4 (right) shows the corresponding trajectories as simulated by the "combined" model when following the (exogenously defined) lead vehicle. Although there are differences, the very simple linear model captures much of real data's dynamics quite well.

It turned out that a simulation of the higher-order model specified in Table 2 does not yield meaningful results: Without non-negativity constraints on the velocities and additional rules for collision avoidance, the trajectories do not make physical sense in that the space ordering of the vehicle is not maintained. An addition of

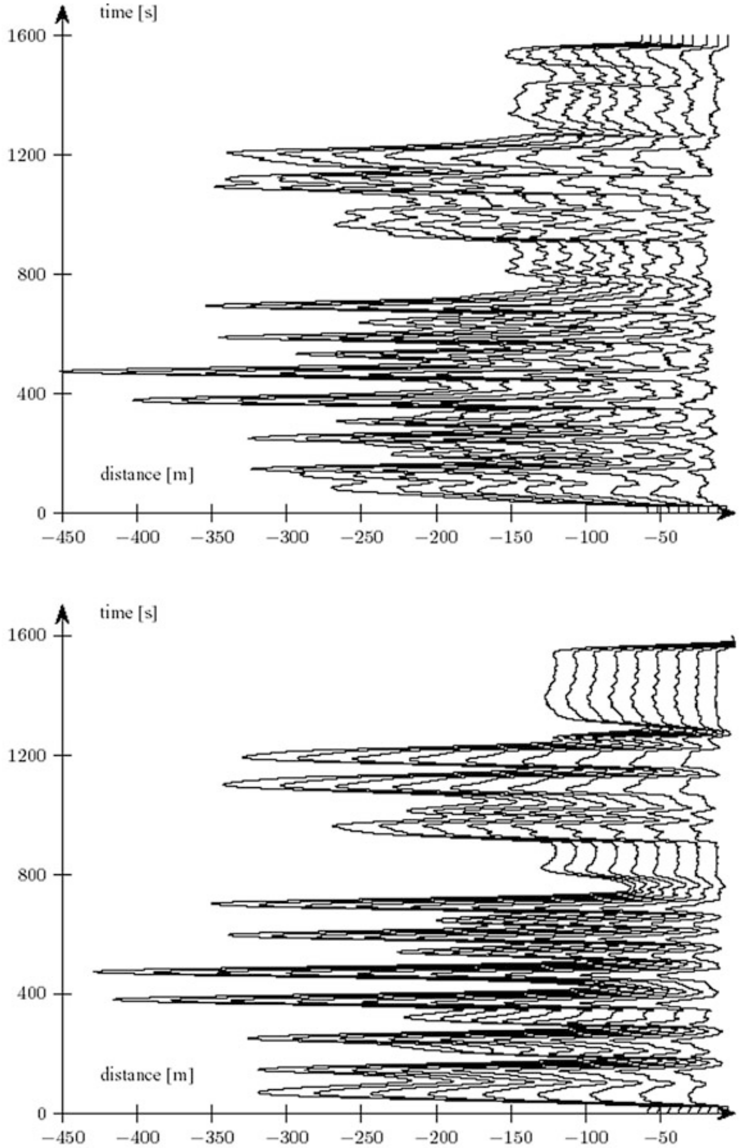

Fig. 4 Space-time plots of original data. X-coordinates are relative to the leading vehicle. *Top*: original data. *Bottom*: the "combined" model

these constraints leads to a configuration where all vehicles follow the lead vehicle with almost zero distance. Apparently, the increased number of parameters of this model enables a better fit, but it also reveals severe structural inconsistencies that most likely result from uncaptured nonlinearities.

3 Summary and Outlook

This study aims at an analysis of how much information can be obtained from real data about a car-following model's parameters. The results obtained from estimating a set of linear models from a rich data set lead to the following observations:

1. A linear model with only three parameters (the "combined" model) already captures much of the dynamics in the data set.
2. The linear model parameters can be estimated robustly and are stable (i) across the given driver population and (ii) across time.
3. Increasing the number of linear model parameters in order to capture higher-order dynamics (the "generalized" model) leads to some increase in fit. This, however, is counterbalanced by implausible simulation results, both without and with bounds on velocities and distances.

As a preliminary conclusion, it may be feasible to state that during most of the time, driver behavior is indeed linear. Further studies should therefore focus more selectively on those episodes in the data where nonlinearities actually take effect.

Acknowledgements We are grateful for the data-set, which has been made available to us by Prof. T Nakatsuji.

References

1. E. Brockfeld, R. D. Kühne, and P. Wagner. Calibration and validation of microscopic traffic flow models. *Transportation Research Records*, 1876:62–70, 2004.
2. G. S. Gurusinghe, T. Nakatsuji, Y. Azuta, P. Ranjitkar, and Y. Tanaboriboon. Multiple car following data using real time kinematic global positioning system. *Transportation Research Records*, 2003.
3. T. Söderström and P. Stoica. Instrumental variable methods for system identification. *Circuits, Systems, and Signal Processing*, 21(1):1–9, 2002.