SUMO User Conference 2026

TIBconferencesession

TIB-OP will set DOI with \TIBdoi

Published: 1970-01-01

Driver Models in SUMO and Elsewhere

Peter Wagner[1]

[1]Institute of Transport Systems, DLR, Rutherfordstr. 2, 12489 Berlin, Germany

[2]Institute of Land- and Sea Traffic, Technische Universität Berlin, Kaiserin-Augusta-Allee 104, 10553 Berlin, Germany

*Correspondence: Peter Wagner, peter.wagner@dlr.de

Abstract: This paper is not strictly scientific. From my point of view, it outlines the major remaining challenges in car-following modelling and highlights some possible future model-based developments of SUMO. It offers a subjective perspective on car following research. Additionally, it is written in a didactic manner, highlighting connections between different research areas without providing conclusive results.

Keywords: Car Following, Intermittent Control, Calibration, SUMO

1 Introduction: modelling drivers

Driver models have been investigated since the 1950s, and even earlier. See, for example, references from this time in [1]–[4]. Even then, it was well known that the complete driving task (which is necessary for creating a solid simulation of a certain area) consists of much more than just car following. Following the taxonomy that has been pointed out in [5] (but not invented there), these are:

- Car Following (CF),
- Lane Changing (LC),
- Intersection Modelling (IM), and
- Routing (which is required for any microscopic simulation involving a network with more than one possible route).

This work here will only investigate CF, as is often the case, thereby paying no heed to Pipes' warning [3]:

> ”The analysis of a line of traffic presented in this discussion is based on postulates that are considerably removed from reality.”

Despite its ”distance to reality”, CF is an important ingredient in traffic flow simulation and has reached a fairly mature state. Nevertheless, research on CF is still ongoing, despite considerable progress having been made since the earliest attempts mentioned above. In more informal terms, it seems that human drivers refuse to be described by any theory and follow their own ideas about how to drive.[1]

As happens in science from time to time, some researchers have believed that the scientific research on a particular topic is complete. This is exemplified by a quote from Richard W. Rothery in chapter 4 of the Traffic Flow Theory Monograph [6] (probably written in 1975, the latest revision is from 1997) who stated (I have shortened his text to improve clarity):

> ”(…) [Car following] is relatively simple compared to other driving tasks, [and] has been successfully described by mathematical models, (…)”

The amazing thing about this quote and Rothery’s text is that he mentions some difficult issues, but does not present them as a different approach. For example, he observes that human control is never correctly described by differential equations (he is one of the authors of one of the most well-known driver models, named by them the L & M model in [7], see also [8]). This will be discussed in section 3.1, and the model itself in section 3.3.

Providing a comprehensive overview of this landscape of driver models is challenging and beyond the scope of this text. However, the work in [9] is an excellent starting point. The authors have undertaken the arduous task of attempting to create a taxonomy of microscopic driving models, which describes the relationships and differences between the various model types. Their paper, with its almost 300 references, is certainly not complete; around 1,000 CF models may be described in the literature. This is in stark contrast to Rothery’s assertion that ”[it] has been successfully described by mathematical models”: if he were be right, we would not require so many models.

1.1 Traffic simulation software

From a practical perspective, simulation software such as SUMO, VISSIM, & others [10], [11] do an amazing job of tackling real transport systems. Many features occasionally appear in theoretical work as new developments, although they are already implemented in existing tools, simply because they are needed to simulate entire transport systems, not just for CF research. The much less glamorous (compared to CF) LC, IM, and routing models are actually needed to simulate real traffic and simply work.

Some examples include emergency braking, road gradients [12], overtaking [13], crash-free driving [14] and, one of my favorites, driver heterogeneity [15]. The latter is well-known to be an important reference for getting the fundamental diagrams right, see also [16] in the SUMO context.

2 Looking into data: an empirical prompt

When it comes to CF, arguably one of the most important representations is that of the CF phase space, as illustrated in Figure 1. This consists of the speed-difference,

[1]Whenever I thought I had made progress and was rewarded with a good match between the simulated and real behaviours, I was back to square one when I switched to a different representation because this other feature does not look realistic.

$\Delta v = V - v$ between the leader's speed, V, and the follower's speed, v, on the x-axis, and the net distance, g (or headway), between them on the y-axis. By definition, it is a phase-space, since the speed difference is the time derivative of the distance g, i.e. $dg/dt = \dot{g} = \Delta v$. All CF data I have seen so far display a similar distribution $p(\Delta v, g)$ in this phase space (see Figure 1). All trajectories display some kind of noisy oscillation, which reminded a former colleague of mine of a broken controler [17]. Note, that the data used are from a 2012 German project named simTD (see [18] for more information about this dataset), but other data I have seen display similar patterns. While the displayed trajectory (orange) is from one car, the distribution behind it (gray colors) is from many cars. However, it is clear that even a single trajectory spans a considerable part of the underlying distribution (especially in terms of Δv), suggesting that a single driver already has a fairly wide distribution in Δv and g. The distribution, but not the trajectories, can even be measured from single car loop detector data.

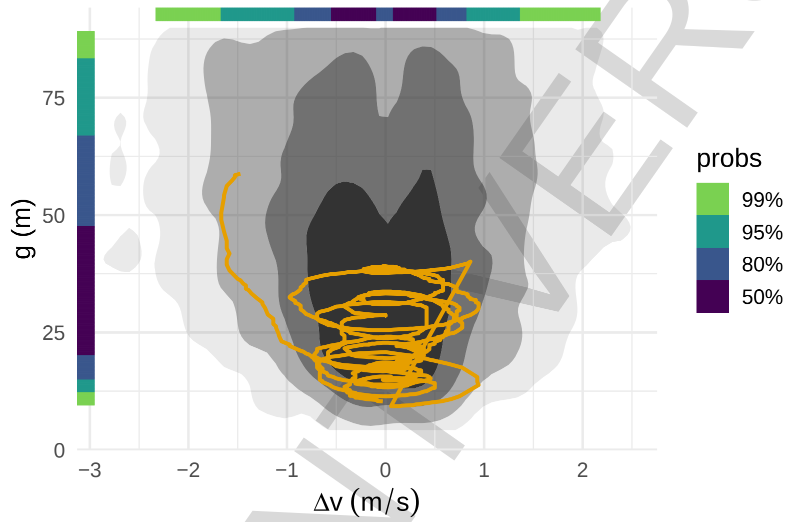

Figure 1. *The grey-shaded areas show the distribution of points $p(\Delta v, g)$ in the CF phase space, and the orange line represents a CF trajectory over 162 seconds, with at least one change in lead vehicle. Additionally, the gap distribution is indicated on the left and the Δv distribution is indicated at the top.*

Another empirical fact is what might be termed the "microscopic fundamental diagram", i.e. a plot of g versus v or vice versa. This is shown in Figure 2, again as a density plot. By cutting along lines of constant speed $v = v_0$, or constant $g = g_0$, respectively, one can sample one-dimensional distributions $p_{v=v_0}(g)$ and $p_{g=g_0}(v)$ where the mode can be extracted by e.g. a kernel density estimate. This yields either the preferred distance as a function of speed (the red points and the fitted red line in Figure 2), or the preferred speed as a function of gap (the green points and the green line). The two lines are not a fit to all points to clarify the effects: for small g, we get $g(v) = g_0 + v\tau$ where $g_0 = 3 \pm 0.3$, $\tau = 0.95 \pm 0.03$, while the green points could also represent a constant speed. Together, these two function constitute what is usually known as the optimal velocity function, if such a thing exists.

The interesting point here is that all the distributions displayed so far are fairly wide. Therefore, one may question the idea that there is such a thing as an optimal velocity function that can be translated into a macroscopic curve for the macroscopic fundamental diagram [19], where speed is drawn as a function of the average inverse gross headway, i.e. the density k. Here, $k = 1/(g + \ell)$, where ℓ is the average length of the vehicles under consideration.

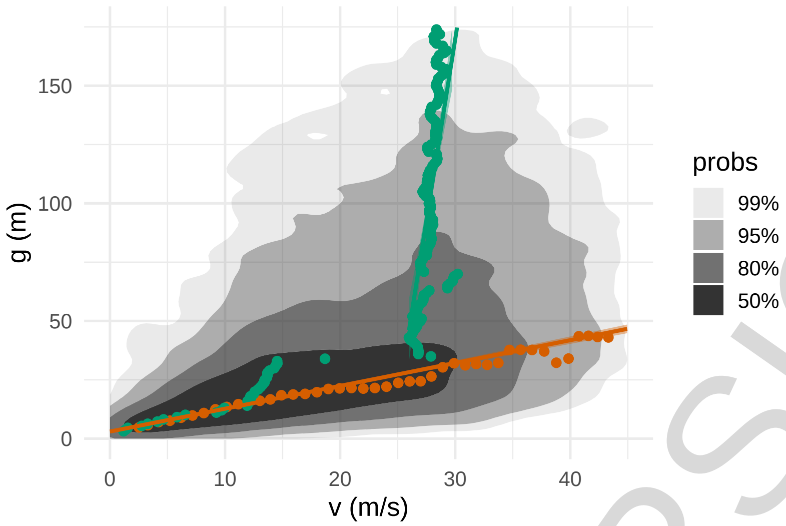

Figure 2. *The distribution (v, g) points, together with points that are drawn at the preferred gap (red) of preferred speed (green). See text for more details.*

A more detailed analysis can be obtained by examining the acceleration. This is shown in Figures 3 and 4, respectively. The analysis there demonstrates that the acceleration is close to the case where $a = 0$ is linear, but there are non-linear effects to be seen in these plots, unfortunately in areas where also the smoothing function's error is large. In all the plots, the smoothing function is a GAM (generalized additive model) which is the default in the tool "R" (see [20] and [21]). Note, however, that the analyses presented in these figures are highly aggregated, so they may miss some details.

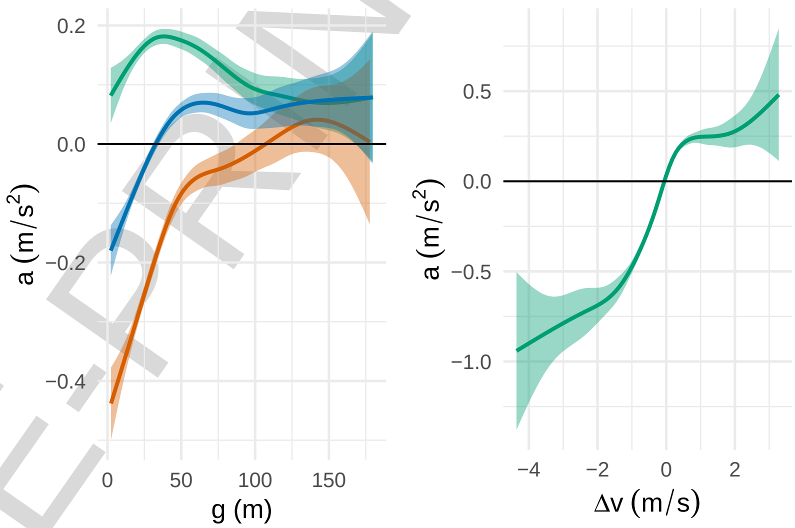

Figure 3. *Acceleration as a function of g (left) and Δv (right). In the left-hand plot, positive values of Δv are shown in green, neutral ones $|\Delta v| < 1\,m/s$ in blue, and negative ones in red. The speeds are picked from $v \in [22, 28]\,m/s$. For the right plot, the data have also been filtered so $g \in [20, 40]\,m$ which is around the zero-crossing of $a(g)$ in the left plot.*

Figure 4 shows the acceleration function in more detail. Data were collected in bins in the $(\Delta v, g, v)$-space, and averaged to obtain the acceleration in the corresponding bin. Typically, and particularly for smaller acceleration values, the standard deviation of the acceleration in a bin is larger than the mean value itself. Only bins containing more than 100 data points are shown here. As can be seen, the acceleration is on average zero near the $\Delta v = 0$ line.

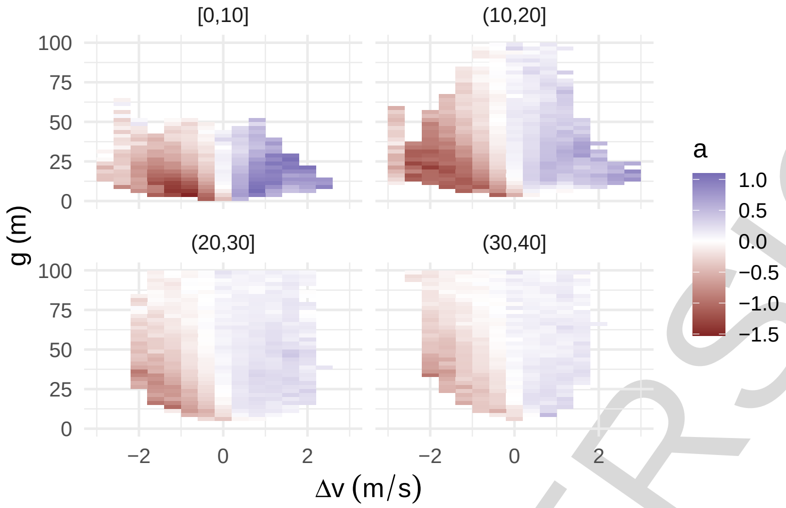

Figure 4. The acceleration function, $a(\Delta v, g)$, is shown for different speeds of following vehicle. Four speed intervals have been selected, with increments of $10\,m/s$, as indicated in the subplot headers.

2.1 A preliminary model consideration

Since the data displays these oscillatory motions, it is important to note that they are not oscillations in the sense of a physics textbook. A real oscillator would avoid the area of the fixed point where the empirical distribution is located, resulting in a completely different ring-shaped distribution. However, it is also important to note that these oscillations are not entirely unexpected. Firstly, there is region in which human drivers feel comfortable following a lead vehicle at a small speed difference, and at an acceptable distance. It should be noted that small Δv are impossible for a human driver to resolve, therefore leading to some ambiguity (see also section 3.1). The same applies to the distance: if it is in an acceptable range, humans are averse to changing it. Most CF **models** have abstracted this to the notion of a stationary point $g^\star = v^\star \tau$, and a few of them have the more general definition of a regime of stationary points $[g_1^\star, g_2^\star]$ where $a = 0$ and $\Delta v = 0$. The value of τ is the driver's preferred time headway and is one of the most important parameters in SUMO's driver models for determining the capacity of a road. These stationary points are solutions to the acceleration function $A(g, v; V) = 0$ on the r.h.s. of the CF differential equation:

$$a = \dot{v} = A(g, v, V, \cdots). \tag{1}$$

Here $a = \dot{v}$ is the vehicle's acceleration. This function can be expanded into a Taylor series in the vicinity of the stationary point. Upon keeping only the first order, another well-known model is obtained: the Helly model [4].

$$a = \frac{1}{T_1}\Delta v + \frac{1}{T_2}\left(\frac{g}{\tau} - v\right). \tag{2}$$

In this equation, the various constants are the derivatives of the acceleration function at the critical point. For example, $1/T_1 = \partial A/\partial \Delta v|_{g=v\tau,\Delta v=0}$ and $1/T_2 = \partial A/\partial g|_{g=v\tau,\Delta v=0}$. These constants are written in this form since they now have a neat descriptive interpretation: T_1 is the relaxation time for the speed difference, and T_2 is the relaxation time for the difference between the stationary speed $v^\star = g^\star/\tau$ and the actual speed.

This model is a fundamental component of older driver assistance systems, and it is implemented in SUMO as the ACC and CACC driver model [22], [23].

This equation can be expressed in the form of a driven harmonic oscillator (with noise, as discussed in section 3.2), as found in physics textbooks:

$$\dot{\Delta v} = -\left(\frac{1}{T_1} + \frac{1}{T_2}\right)\Delta v - \frac{1}{T_2\tau}g + \dot{V}(t) + \frac{1}{T_2}V(t), \tag{3}$$

$$\dot{g} = \Delta v. \tag{4}$$

The change to this form also changes the coefficients. Although it is not immediately obvious, setting $\dot{\Delta v} = 0$ and $\Delta v = 0$ still gives the correct stationary point $g^* = V\tau = v^*\tau$, since $v^* = V$. Depending on the figures for the three time constants, such a model can display damped oscillations and even resonance, provided that the driving force is oscillatory. To my knowledge, this has never been used in CF research. The driving force is the combination of the leader's speed and acceleration $F(t) = \dot{V}(t) + V(t)/T_2$.

Two things should be noted here. Firstly, in the text above, the word "model" is highlighted for a reason: in my view, it is unclear whether the concept of a stationary point is applicable in reality. The Wiedemann model [24] has the idea of a no-action zone near $\Delta v = 0$ (where a remains constant from when the driver enters the no-action zone), and one might question whether drivers spend a significant amount of time with $a = 0$. However, the acceleration distribution is strongly peaked at $a = 0$, lending some credibility to the fixed point idea. The construction in section 3.1 gives an example of a model that does not require an actual fixed point or a line of fixed points.

Secondly, SUMO's default model (affectionately named the SK model from its creator, Stefan Krauß) cannot reproduce this distribution without additional adjustments. To achieve the correct width of the phase space distributions $p(\Delta v, g)$, one requires either a fleet of vehicles or the driver state device (DSD) [25], or one of the other models in SUMO's driver model zoo to get the correct width of the distributions. See Figure 5 for a comparison of the default model, the default model with the driver device and the empirical data from Figure 1 (the 95% line).

It is still an open question, whether the width of the distributions really requires distributed parameters, or whether it can already be produced by a carefully chosen model. For example, it could be argued that the time headway of each driver, denoted by τ, is not fixed, but rather follows a stochastic equation. This would result in a wide range of time headway values, and subsequently, headway values (see [26]). However, while this is a reasonable model, it introduces two additional parameters and a further mechanism that are difficult or even impossible to access empirically.

3 Challenges of driver models (a subjective selection)

3.1 Intermittent control

Human control is never continuous in either time or state (e.g. speed control). This will be demonstrated by having a closer look at the model in eq. (3). This time, we will assume a lead vehicle travelling with constant speed V:

$$\dot{\Delta v} = -\left(\frac{1}{T_1} + \frac{1}{T_2}\right)\Delta v - \frac{1}{T_2\tau}g + \frac{1}{T_2}V(t). \tag{5}$$

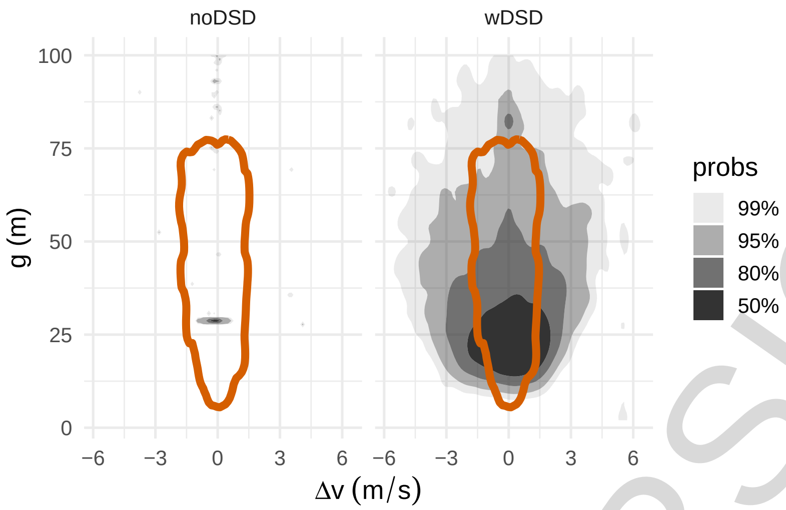

Figure 5. *A comparison of the phase space distribution of SUMO's default model (left), with the same model equipped with SUMO's driver state device (right). The 95% line from the empirical data in Figure 1 is overlaid in orange.*

Formulating it as a differential equation (ODE) assumes that humans control acceleration at any given millisecond (i.e. continuous control!), even for very small changes in Δv. Clearly, these assumptions are unrealistic, and can only be justified by testing to ensure they do not produce unacceptable results[2]. In fact, the assumptions behind Helly's model are very reasonable and are in line with a simple driver model, in which human controllers simply balance the two deviations from their preferred state. Adding an intermittent control mechanism (it is often named action points, AP), as stated by [27] and acknowledged by Rothery [6], changes the behaviour of the model considerably. At first glance, this looks similar to SUMO's DSD, but it is quite different. It also differs from the psycho-physical models that began with [27] and continued in the modelling behind VISSIM, which was introduced by Wiedemann [24]. See also [28] for a modern approach to intermittent control and [29] for similarities with human balancing.

These psycho-physical models, as they are known, define thresholds in phase space where the driver's behaviour changes, for example from approaching to following. We believe that these thresholds are far too static. Therefore, we assume that the drivers monitor speed difference, distance, speed, and maybe, acceleration, and only act if they recognize significant changes in these variables as a reason to do so. This can be incorporated into an intermittent control model (named Helly-AP) as follows:

$$\Delta v(t + \Delta t) = \Delta v(t) + \Delta t\, a(t) \tag{6}$$

$$a(t) = \begin{cases} a_i, & \text{if } |\Delta v(t - t_i) - \Delta v(t)| < \theta_1 \\ & \text{and } |g(t - t_i) - g(t)| < \theta_2 \\ -\left(\frac{1}{T_1} + \frac{1}{T_2}\right)\Delta v - \frac{1}{T_2 \tau} g + \frac{1}{T_2} V(t), & \text{else} \end{cases} \tag{7}$$

Here, the variable t_i represents the last time acceleration changed, and a_i represents the acceleration computed at that time. The driver retains her or his last acceleration until something changes – in this case, until a considerable change in Δv or in g occurs.

[2] One may note that even ADAS controllers cannot achieve this; they have cycle times, hopefully in the millisecond range or below, and their measurement of Δv etc. are noisy. This at least prevents them from responding to arbitrary small changes in Δv since then they would follow the noise.

This model is compared in Figure 6 to a standard Helly model with the same parameters and the same initial condition. Note, that both threshold conditions (for Δv and g) are needed: if the driver controls only Δv, the intermittent dynamics may lead to crashes or ever increasing distances, since control in Δv finally reaches a corridor where the vehicle bounces between the thresholds $\Delta v = \pm\theta_1$, but distance does not settle.

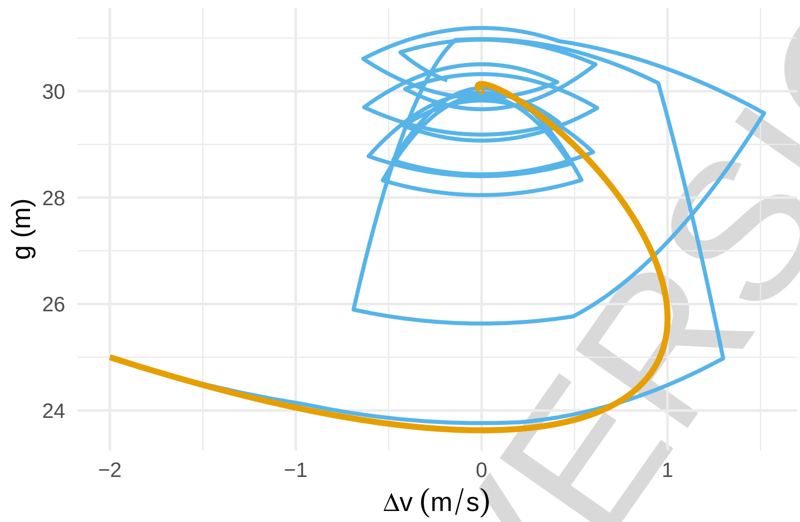

Figure 6. *Two simulated trajectories are shown: one following the traditional Helly model (orange), and one following the Helly-AP variant with intermittent control (blue).*

As can be seen from Figure 6, the intermittent model does not settle into a fixed point. Although the two trajectories originally look similar, they exhibit very different behaviour. The width of the Helly-AP oscillatory behaviour is controlled by the two thresholds, and can be made smaller or larger, but it does not vanish. Note, that the models in both cases are deterministic. However, the behaviour of the Helly-AP model bears some hallmarks of deterministic chaos, since a small deviation in the initial conditions develops into a large deviation after a few APs have been performed by the model[3]. Clearly, both approaches bear little resemblance to the observed phase space density, but this can be addressed by implementing other features, and/or adding mild stochasticity. Assuming fixed thresholds does not reflect how human drivers act. Therefore, the thresholds may change after each activation, and they depend on the driver's state.

Intermittent control is more complicated to model than continuous control, which may explain why it has not found many applications in reality or in theory. The same argument applies as for a model with fluctuating τ: another mechanism is added, which in this case requires two additional parameters that are difficult to measure empirically. Another reason for the neglect of this approach was stated by Rothery:

> ”Inertia, on the other hand, both in the operator and the machine, creates an appearance of smoothness and continuity to the control element.”

As one might add, measurement noise can also mask action points, but they can be uncovered from the data [30], at least when the acceleration jump is large enough to stand out from the measurement noise.

Note that, although this approach is similar to SUMO's DSD, it is different. The DSD initiates an update to the acceleration if the difference between the perceived $(\tilde{\Delta v}, \tilde{g})$ and

[3]As an aside, one might note the similarity to what in physics is named a billiard.

the real $(\Delta v, g)$ is larger than a certain threshold. In addition, the thresholds depend on distance, which is of course a reasonable assumption.

An interesting observation based on simulations with this model is that, while the normal Helly model is string unstable over a wide range of parameters – a problem that also effects driver assistant models (see [31]) – the Helly-AP model in no longer string unstable. However, instability is another thorny issue: not all congested situations where string instability may play a role result in a traffic jam, so there is no clear connection between string instability and traffic jams. Additionally, a simple addition to the Helly model or the model in eq. (8) may alter the stability slightly, in the form of a cap on acceleration. The Helly model, as stated in eq. (2) has no acceleration bound, which does not change the start of the oscillation, but it may reduce its effect.

In summary, a model with intermittent control requires more detailed modelling, making it a more complicated model. It adds two more parameters to the three parameters of Helly's model: the two thresholds θ_1 and θ_2. Additionally, the update mechanism must remember the variables $a_i, \Delta v_i$, and g_i, i.e. the values at which the last action occurred. Furthermore, it is not immediately clear how to describe stability, which is easily achieved in continuous control (see [32]). Further research in this area is clearly needed, and we believe it will be a worthwhile endeavour.

3.2 Modelling stochasticity

Car following models, especially from physicists, often describe stochasticity by adding white noise to the acceleration. While this approach has been used successfully to describe random walks, it is inadequate when applied to vehicles: these are heavy objects, so their stochasticity cannot be modelled by assuming changes in the acceleration within milliseconds. Therefore, in my view, models that do so should not to be taken too seriously, although they can still be useful. SUMO's default model being the prime example of this. This could be remedied by either adding coloured noise which is a smoothed variant that avoids such large jumps, or by using the intermittent control mentioned above. Noise can enter at three points: in the perception of the distances, speed differences, and speeds (see the DSD again), in the intermittent control thresholds, or at the point where the new acceleration is computed. In this case, there is no problem with too large jumps in acceleration caused by white noise, the driver simply applies a new acceleration which may introduce an additional actuation error with any distribution, provided it is bounded.

This can be applied to the Helly-AP model. Unfortunately, this does not simplify modelling, since at least one additional step is needed, with another parameter: the acceleration noise. Even this particular model, which started out so simply, has become a fairly complicated construct, and one might feel, that some of the parameters are redundant. Addressing this is an interesting avenue for future research. Nevertheless, a quick simulation of a platoon of those Helly-AP cars with stochastic thresholds and in the acceleration noise shows some promising results, as can be seen in Figure 7.

3.3 Reaction times

Humans have reaction times r reaching from $r = 0.2\,s$ for world-class 100 metre runners in the starting blocks to several seconds when a conscious decision has to be made.

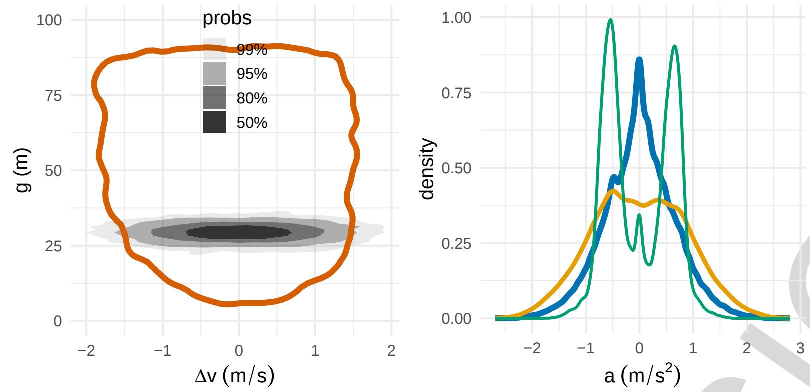

Figure 7. *This is a simulation of a platoon of 21 vehicles of the stochastic Helly-AP model. On the left: While the model displays little variation in g, the variation in Δv seems close to the empirical results (orange line from Figure 5). The plot on the right shows the acceleration distribution for the deterministic model (green), the probabilistic model (blue, thick), and the acceleration in the probabilistic model at the action-points (orange).*

Clearly, humans also have foresight, which is used to counteract the effect of reaction times. A prime example of this is when we catch a ball, a process whereby the compute machinery in our brain simply takes care of all the delays in the actuator loops in our body. The same holds for CF, which is why I have never been a fan of delay differential equations (DDE) in CF[4]. Of course, there is the notorious L & M model:

$$a = \frac{1}{T_1}\Delta v(t-r)\,\frac{v^m}{g(t-r)^\ell}, \tag{8}$$

which includes reaction time and was one of the first models to propose a mechanism for traffic instabilities. These instabilities stem from the finite human reaction times, and/or the string instability in a platoon of drivers if the model and the parameters are chosen unwisely. However, we now know that reaction times are not necessary to produce unstable platoons, see [32] for an excellent overview on this topic (and other topics as well). Despite its popularity, the model eq. (8) is not a good model: for small g, the acceleration explodes. If there were no reaction time, this model would be crash-free. However, for small g, the acceleration changes from a large negative value to a large positive value with very small changes of Δv. Additionally, the acceleration is unbounded. In fact, this is necessary for the oscillations to be unbounded for very large r: if the acceleration is limited, the oscillation amplitude no longer increases over time but remains limited. Finally, the model has a line of fixed points for $\Delta v = 0$, without any preferred distance. This is in contrast to the empirical data, and should therefore be an important ingredient of any reasonable CF model. See also Figure 3 for an idea about the relationships displayed by the data between a and g and a and Δv, to be compared with the ones assumed by the L & M model[5].

Coming back to the reaction times, I have observed instances in data where reaction times are negative. This occurs when a driver approaches a traffic signal and begins braking before the lead vehicle, creating the appearance of a negative reaction time. See, e.g., Figure 2 in [33]. Clearly, this can also happen in other circumstances if drivers see the braking lights of the next-nearest car, causing them to start braking earlier than

[4] During my PhD, I actually worked with DDEs, so I am familiar with them and find them a fascinating topic.
[5] This can be easily fixed by exchanging Δv by $\Delta v + (g/\tau - v)$ in the L & M equation. However, I cannot fathom their reasoning behind choosing v^m, since acceleration should decrease with increasing speed. Consequently, direct fits of this model to data yield a negative m. In their own work they often choose $m = 0$.

the lead car. Of course, not all drivers do this, and this is precisely the challenge with reaction times: they are strongly context-dependent. In normal CF situations, I think they are masked by the same mechanism that allows us to catch a ball: we anticipate what will happen and act accordingly. It only fails if drivers are not paying attention, or if something unexpected happens such as an unanticipated lane-change or a harsh braking manoeuvre by the vehicle in front due to a stimulus we have not seen either.

Note, that the Helly model can be interpreted differently depending on how it is written:

$$a = \alpha(g + \Delta v T_a - v\tau)$$

where the term $g + \Delta v T_a$ can be interpreted as an anticipated future gap $g(t + T_a) := g + \Delta v T_a$. To align more closely with the reaction time formulation, one could write a model with an explicit anticipation mechanism by assuming that the driver does not use $g(t - r)$ but rather the anticipated value $g(t - r) \rightarrow \hat{g} = g(t - r) + r\Delta v(t - r)$:

$$a = \left(\frac{1}{T_1} + \frac{r}{\tau T_2}\right)\Delta v(t - r) + \frac{1}{T_2}\left(\frac{g(t - r)}{\tau} - v(t - r)\right). \tag{9}$$

However, in this form this has only a weak effect on stability, since the prefactor of Δv gets slightly bigger.

The final challenge with reaction times is that they are difficult to obtain empirically, as they are often confused with relaxation times. I have tried this several times without achieving a satisfactory result. The main idea is to calculate relaxation times using the autocorrelation function of the acceleration or speed difference, which, on average yields results between $T_{\text{rel}} = 2 \ldots 5\,s$. To get a hold on reaction times, one should relate the two speed-time series of the leader $V(t)$ and the follower $v(t)$, and find the delay between them where a the mismatch $\Delta v_r = |V(t) - v(t - r)|$ is minimal. This should be done as a function of time as discussed above, since reaction times are most likely time and state dependent. However, there may be better approaches, we have experimented a bit with dynamic time warping [34], but so far without decisive results.

3.4 Calibration challenges

Considerable work has gone into calibrating CF and SUMO. Calibration is essential when applying a simulation model to real-world problems. This can be done with good accuracy, see, for example, Yun-Pang Flötteröd's work [35] with open data from DLR's research intersection [36]. However, we learned that one should be careful: in free traffic, it does not make much sense to calibrate the preferred headway τ, or the distance at standstill, g_0. Similarly, it is of no use to calibrate maximum speeds in dense traffic, since the vehicles do not come anywhere near this maximum speed.

Another challenging situation arises when we have the wrong model, which is often the norm rather than the exception. I worked on this topic for a while: the assumption was, that the parameters of a linear model like Helly's model are not constant, but change over time [37]. However, I learned that such an observation may also result from not having the correct model. For example, running a simulation with SUMO's default model and then using a linear approximation to this model makes the parameters of the linear model time- or state-dependent. From this perspective it is impossible to determine whether these parameters are genuinely time-dependent or if we simply have the incorrect model.

The unpublished work of my colleagues in the SUMO group points to other indirect effects when doing calibration. This can be exemplified as follows: The model under consideration is the IDM [38], which, in the absence of a leading vehicle has a very simple acceleration function which can be described by just three parameters, the maximum speed (traditionally named v_0 in IDM), the maximum acceleration a_0 and the exponent δ. Typically, for this parameter the value $\delta = 4$ is used. In this special case, the model can be simplified further by using scaled variables $\tilde{a} = a/a_0$ and $\tilde{v} = v/v_0$, this leads to $v_0 = 1$:

$$\tilde{a} = 1 - \tilde{v}^{\delta} \tag{10}$$

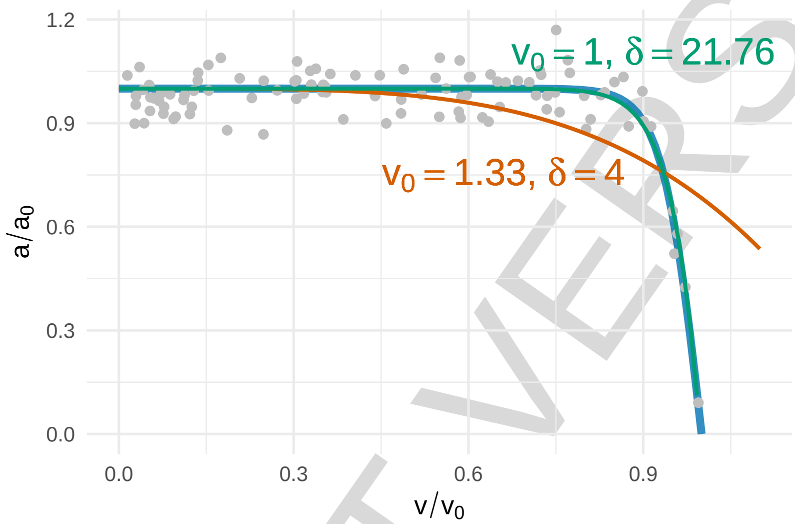

Figure 8. *A synthetic dataset constructed with a different acceleration function than the one of the IDM, in this case $a(v) = tanh(-15(v-1))$ was used with some noise added. When fitting IDM's v_0 by keeping $\delta = 4$ fixed, the fit fails (red line), consistent with what has been observed with real data. Allowing the fitting of both v_0 and δ (green line) results in a good fit, albeit with a very large δ.*

An interesting observation was that in order to fit this to real data, the maximum speed was completely wrong; it was much larger than what had actually been observed in the dataset used for the calibration. After some experimentation, the reason for this mismatch became apparent, see Figure 8 for a visualization. Real data may not follow eq. (10), they display a relatively sharp cut-off near the maximum speed. However, with the choice $\delta = 4$, the decline in a is much smoother than in the data. As long as one kept δ fixed during calibration, the optimizer had no choice but to increase v_0 beyond what was observed in the data. Therefore, calibrating both δ and v_0 fixed the issue, but increased δ to the fairly large value of roughly 20.

While the mechanisms leading to calibration failure are transparent in the cases mentioned here, this may be overlooked in other cases. Furthermore, part of the perceived failure stems from the fact that one has a mental picture, especially of the parameter v_0, which is at odds with the initial calibration result and the data. If we obtained a result for the parameters (ℓ, m) in the L & M model, no one would be surprised, because we have no intuition about those parameters except that the exponent of g is greater than zero. Even the poor fit to the IDM might be acceptable until it is used to extrapolate to scenarios, where higher speeds are needed than the ones in the calibration. Apart from the usual advice to be careful and to look at one's results from as many different angles as possible, I do not have a final solution to these kind of challenges.

4 Final remarks

In my view, the challenges mentioned in section 3 constitute an interesting research program. I would like to add two more that need to be addressed, particularly for inclusion into practical work:

- In my view, the most exciting aspect of driver modelling is enabling the models to realistically simulate crashes.
- Modelling bicycle drivers [39].

We have already discussed some areas where simulation models could be improved or applied to real-world problems. We recently started this endeavour by adding a blackout mechanism to SUMO (which is still only in Python, not in the main repository). This can cover at least the general weekly trends in crash numbers, since these are strongly correlated with the traffic flow. We are currently working on the finer details, such as the distribution of different crash types that may depend on the state of traffic: for example, there are many more single-vehicle accidents at night or during periods of low traffic than during busy times. For more information, please refer to Ronald Nippold's work in this volume.

Finally, there is the modelling of cyclists. In SUMO, this uses the default model of SUMO with a re-parameterisation. This is not completely wrong, but clearly does not take the specific characteristics of cyclists into account. For starters, real cyclists do not adhere to such a strict no-crash approach as in the default CF model. This is because they believe that they can swerve around the lead cyclist easily in case something unexpected happens. Also, in terms of capacity, cyclists do not ride single file, even on one-lane cycle paths, which could increase the actual capacity beyond the expected level. This could be remedied using SUMO's in-house tools, for instance by either using the sublane model on bicycle paths with at least two sublanes or giving each bicycle path two lanes by default. However, this may overestimate capacity, since drivers may be reluctant to overtake another cyclist by using part of the pedestrian sidewalks. Unfortunately, we do not have much data on this.

Data availability statement

The plots in this text had been done by a quarto script, which can be obtained upon reasonable request from the author, together with the data needed to make them.

Author contributions

The text is written entirely by Peter Wagner. In line with CRediT, I have done the Conceptualization, the data analysis, the simulation, the writing and partly the review. I had lots of help from an LLM regarding language, a little for scanning the literature. Many thanks to Robert Alms for the final review, pointing out some glitches and improving language. All remaining errors are of course my own.

Competing interests

Apart from the fact that I am a big fan of SUMO for obvious and transparent reasons, I declare that I do not have to state any competing interests.

Funding

Most of the funding for my research comes from the German Aerospace Center's (DLR) basic funding and third-party funding from a variety of projects, most of which are funded by German sources such as the Ministries of Transport, Economics, and Science.

Acknowledgments

It is a pleasure to thank everyone who has worked with me over the last 30 years (and perhaps endured me, hopefully not too much). It is impossible to name everyone, especially the many people in the community with whom I have discussed things, so I hope this general thank you will suffice. As this text shows, I stuck to the basics for the most part; more could have been done. Finally, I would like to thank my employers at the DLR and, at the beginning, the Center of Parallel Computing in Cologne, for providing the means to carry out this work.

References

[1] E. M. J. Herrey and H. Herrey, "Principles of Physics Applied to Traffic Movements and Road Conditions," *American Journal of Physics*, vol. 13, no. 1, pp. 1–14, Feb. 1945, ISSN: 0002-9505. DOI: 10.1119/1.1990647.

[2] A. Reuschel, "Fahrzeugbewegung in der Kolonne bei gleichförmig beschleunigtem oder verzögertem Leitfahrzeug," *Zeitschrift des österreichischen Ingenieur und Architekten-vereins*, p. 95, Jul. 1950, In German.

[3] L. A. Pipes, "An Operational Analysis of Traffic Dynamics," *Journal of Applied Physics*, vol. 24, no. 3, pp. 274–281, Mar. 1953. DOI: 10.1063/1.1721265.

[4] W. Helly, "Simulation of bottlenecks in single lane traffic flow," in *Proceedings of the symposium on theory of traffic flow.*, 1959, pp. 207–238.

[5] J. Erdmann, personal communication, May 25, 2025.

[6] N. H. Gartner, C. J. Messer, and A. Rathi, Eds., *Traffic flow theory*. TFHRC, 1997, https://www.fhwa.dot.gov/publications/research/operations/tft/index.cfm.

[7] D. C. Gazis, R. Herman, and R. W. Rothery, "Nonlinear follow-the-leader models of traffic flow," *Operations Research*, vol. 9, no. 4, pp. 545–567, 1961. DOI: 10.1287/opre.9.4.545.

[8] D. C. Gazis, "The origins of traffic theory," *Operations Research*, vol. 50, no. 1, pp. 69–77, 2002. [Online]. Available: http://www.jstor.org/stable/3088451.

[9] T. Zhang, P. Jin, S. McQuade, A. Bayen, and B. Piccoli, "Car-following models: A multidisciplinary review," *IEEE Transactions on Intelligent Vehicles*, vol. PP, pp. 1–26, Jan. 2024. DOI: 10.1109/TIV.2024.3409468.

[10] E. A. Algherbal and N. T. Ratrout, "A comparative analysis of currently used microscopic, macroscopic, and mesoscopic traffic simulation software," *Transportation Research Procedia*, vol. 84, pp. 495–503, 2025, Smart Mobility and Logistics Ecosystems, ISSN: 2352-1465. DOI: https://doi.org/10.1016/j.trpro.2025.03.101.

[11] P. A. Lopez, M. Behrisch, L. Bieker-Walz, *et al.*, "Microscopic traffic simulation using sumo," in *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, 2018, pp. 2575–2582. DOI: 10.1109/ITSC.2018.8569938.

[12] T. Ersal, I. Kolmanovsky, N. Masoud, *et al.*, "Connected and automated road vehicles: State of the art and future challenges," *Vehicle System Dynamics*, vol. 58, no. 5, pp. 672–704, 2020. DOI: 10.1080/00423114.2020.1741652.

[13] R. Zhang, B. Qiong, K. Brijs, E. Hermans, Q. Qu, and Y. Shen, "Overtaking maneuvers on two-lane highways under the microscope: Exploration of a multidimensional framework for the analysis of safety, comfort and efficiency using simulator data," *Accident Analysis & Prevention*, vol. 202, p. 107 613, 2024, ISSN: 0001-4575. DOI: 10.1016/j.aap.2024.107613.

[14] D. Liu, B. Besselink, S. Baldi, W. Yu, and H. L. Trentelman, "On structural and safety properties of head-to-tail string stability in mixed platoons," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 6, pp. 6614–6626, 2023. DOI: 10.1109/TITS.2022.3151929.

[15] M. Makridis, A. Anesiadou, K. Mattas, G. Fontaras, and B. Ciuffo, "Characterising driver heterogeneity within stochastic traffic simulation," *TRANSPORTMETRICA B-TRANSPORT DYNAMICS*, vol. 11, no. 1, pp. 725–743, 2023, ISSN: 2168-0566. DOI: 10.1080/21680566.2022.2125458.

[16] J. Erdmann, J. Rummel, and P. Wagner, "Sumo and the german handbook for the dimensioning of highways (hbs)," in *Towards Simulation for Autonomous Mobility*, ser. Lecture Notes in Mobility, Springer, May 2017. [Online]. Available: https://elib.dlr.de/112657/.

[17] J. Rataj, personal communication, May 2010.

[18] P. Wagner, R. Nippold, S. Gabloner, and M. Margreiter, "Analyzing human driving data an approach motivated by data science methods," *Chaos, Solitons & Fractals*, vol. 90, pp. 37–45, 2016, Challenges in Data Science, ISSN: 0960-0779. DOI: https://doi.org/10.1016/j.chaos.2016.02.008.

[19] B. S. Kerner, *Introduction to Modern Traffic Flow Theory and Control: The Long Road to Three-Phase Traffic Theory*. Springer Berlin Heidelberg, 2009, ISBN: 9783642026058. DOI: 10.1007/978-3-642-02605-8.

[20] R Core Team, *R: A language and environment for statistical computing*, R Foundation for Statistical Computing, Vienna, Austria, 2025. [Online]. Available: https://www.R-project.org/.

[21] H. Wickham, *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York, 2016, ISBN: 978-3-319-24277-4. [Online]. Available: https://ggplot2.tidyverse.org.

[22] The SUMO Team. "Sumo driver state device." (2026), [Online]. Available: https://sumo.dlr.de/docs/Car-Following-Models/ACC.html,https://sumo.dlr.de/docs/Car-Following-Models/CACC.html (visited on 02/10/2026).

[23] K. N. Porfyri, E. Mintsis, and E. Mitsakis, "Assessment of ACC and CACC systems using SUMO," in *SUMO 2018- Simulating Autonomous and Intermodal Transport Systems*, E. Wießner, L. Lücken, R. Hilbrich, *et al.*, Eds., ser. EPiC Series in Engineering, vol. 2, EasyChair, 2018, pp. 82–93. DOI: 10.29007/r343. [Online]. Available: /publications/paper/HvLp.

[24] R. Wiedemann, "Simulation des Straßenverkehrsflußes," Institut für Verkehrswesen, Universität Karlsruhe, Tech. Rep., 1974, Heft 8 der Schriftenreihe des IfV, in German.

[25] The SUMO Team. "Sumo driver state device." (2026), [Online]. Available: https://sumo.dlr.de/docs/Driver_State.html (visited on 02/10/2026).

[26] P. Wagner, "Analyzing fluctuations in car-following," *Transportation Research Part B: Methodological*, vol. 46, no. 10, pp. 1384–1392, 2012, ISSN: 0191-2615. DOI: 10.1016/j.trb.2012.06.007.

[27] E. P. Todosiev and L. C. Barbosa, "A proposed model for the driver-vehicle-system," *Traffic Engineering*, vol. 34, pp. 17–20, 1963.

[28] U. Durrani and C. Lee, "A new car-following model with incorporation of Markkula's framework of sensorimotor control in sustained motion tasks," *Transportation Research Part B: Methodological*, vol. 184, p. 102 969, Jun. 1, 2024, ISSN: 0191-2615. DOI: 10.1016/j.trb.2024.102969.

[29] T. Insperger, G. Stépán, and J. Milton, "Dynamics of human balancing," in *Controlling Delayed Dynamics: Advances in Theory, Methods and Applications*, D. Breda, Ed., Cham: Springer International Publishing, 2023, pp. 343–364, ISBN: 978-3-031-01129-0. DOI: 10.1007/978-3-031-01129-0_11.

[30] P. Wagner, J. Erdmann, and R. Nippold, "Action-points in human driving and in sumo," *SUMO Conference Proceedings*, vol. 1, pp. 113–122, Jul. 2022. DOI: 10.52825/scp.v1i.104.

[31] B. Ciuffo, K. Mattas, M. Makridis, *et al.*, "Requiem on the positive effects of commercial adaptive cruise control on motorway traffic and recommendations for future automated driving systems," *Transportation Research Part C: Emerging Technologies*, vol. 130, p. 103 305, 2021, ISSN: 0968-090X. DOI: https://doi.org/10.1016/j.trc.2021.103305.

[32] M. Treiber and A. Kesting, *Traffic Flow Dynamics: Data, Models and Simulation*. Springer Nature Switzerland, 2025, ISBN: 9783031939228. DOI: 10.1007/978-3-031-93922-8.

[33] P. Wagner, "How human drivers control their vehicle," *The European Physical Journal B*, vol. 52, no. 3, pp. 427–431, Jul. 2006, ISSN: 1434-6036. DOI: 10.1140/epjb/e2006-00300-1.

[34] P. Tormene, T. Giorgino, S. Quaglini, and M. Stefanelli, "Matching incomplete time series with dynamic time warping: An algorithm and an application to post-stroke rehabilitation," *Artificial Intelligence in Medicine*, vol. 45, no. 1, pp. 11–34, 2008. DOI: 10.1016/j.artmed.2008.11.007.

[35] Y.-P. Flötteröd and P. Wagner, "Leveraging intersection trajectory data for microscopic traffic simulation," in *Transport Systems Development – Methods and Solutions*, G. Sierpiński, S. Naumann, and E. Macioszek, Eds., Cham: Springer Nature Switzerland, 2026, pp. 12–24, ISBN: 978-3-032-14826-1.

[36] C. Schicktanz, L. Klitzke, K. Gimm, *et al.*, *DLR Urban Traffic dataset (DLR UT)*, version 1.3.0, Zenodo, Jul. 2025. DOI: 10.5281/zenodo.15754836.

[37] P. Wagner, G. Flötteröd, R. Nippold, and Y.-P. Wang, "Simplified car-following models," in *Proceedings of the Transportation Research Board*, 91, Transportation Research Board, Jan. 2012, pp. 1–11.

[38] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Physical Review E*, vol. 62, pp. 1805–1824, 2 Aug. 2000. DOI: 10.1103/PhysRevE.62.1805.

[39] A. Roosta, H. Kaths, M. Barthauer, J. Erdmann, Y.-P. Flötteröd, and M. Behrisch, "The state of bicycle modeling in sumo," *SUMO Conference Proceedings*, vol. 4, pp. 55–64, Jun. 2023. DOI: 10.52825/scp.v4i.215.