Adaptive Driver Reaction Parameters in the Optimal Velocity Model

Pavlos Constas-Malevanets   Yang Yang Zhang

December 19, 2025

Abstract

Stop-and-go waves, or phantom traffic jams, emerge spontaneously in uniform traffic flow without external bottlenecks. The classical Optimal Velocity Model (OVM) captures the onset of such instabilities but fails to sustain them: oscillations decay as the system relaxes toward uniform flow. We propose the Delayed Adaptive OVM (DA-OVM), which replaces the constant driver sensitivity a with a state-dependent function $a(h, \Delta v)$ that increases when headway is small or the gap is closing, combined with an explicit reaction delay τ. Fitting six car-following models to empirical ring-road data from Sugiyama et al., we find that the DA-OVM achieves the lowest RMSE (0.865 vs. 0.920 m/s) and highest R^2 (0.181 vs 0.0735), outperforming the baseline OVM. Crucially, the DA-OVM sustains velocity oscillations at 63% of the observed amplitude, compared to 52% for constant-sensitivity models. The fitted parameters reveal that drivers react nearly three times faster when gaps are small ($\beta_h = 1.77$) and exhibit strong sensitivity to closing speeds ($\beta_v = 0.80$ s/m). These results demonstrate that adaptive sensitivity and delayed feedback together capture the persistent, periodic structure of phantom traffic jams better than the classical OVM.

1 Introduction

Stop–and–go waves in traffic, often called *phantom traffic jams*, are nonlinear density waves that form even when all cars move on a uniform road with no bottlenecks, lane changes, or external disruptions. Unlike congestion caused by physical obstacles, phantom jams emerge purely from driver interaction dynamics. These self-induced waves propagate backward relative to the direction of motion, despite all drivers intending to move forward, and result from delayed or imperfect velocity corrections that amplify small perturbations.

The *Optimal Velocity Model* (OVM), introduced by Bando *et al.* [1], presents a minimal system of coupled ordinary differential equations describing how drivers continuously adjust their velocity toward a desired, headway-dependent target speed. The model is conceptually simple yet captures fundamental features of traffic instability on closed rings and highways. In the classical OVM, the driver sensitivity parameter $a > 0$ is a constant for the entire system, which assumes that all drivers correct their speed at the same rate, regardless of conditions.

However, real driver behavior contradicts this assumption. Human drivers accelerate and brake with varying urgency: when the gap h to the leading vehicle is small, drivers react aggressively, whereas for larger gaps, corrections are slower and smoother. Drivers also compensate asymmetrically to relative velocity $\Delta v = v_{i+1} - v_i$, reacting most strongly when closing in on a slower predecessor ($\Delta v < 0$) [2]. A constant a cannot capture these adaptive effects, and choosing it poorly leads to less realistic results.

Furthermore, while the classical OVM can initiate stop-and-go waves from small perturbations, it fails to sustain them over extended time horizons. In ring-road experiments [3], traffic jams persist indefinitely, with vehicles repeatedly cycling through acceleration and braking phases. The OVM, by contrast, exhibits overdamped behavior: after an initial transient, velocity fluctuations decay and the system relaxes toward uniform flow. This occurs because the constant sensitivity a acts as a uniform restoring force that smooths out disturbances regardless of their magnitude or the local traffic state. The model captures the onset of congestion but not its persistence, a critical limitation when studying the periodic structure of real phantom jams.

In this paper, we propose to replace the constant a with a state-dependent function $a(h, \Delta v)$ that modulates driver responsiveness based on instantaneous traffic conditions. We further incorporate an explicit reaction delay τ to model the finite time required for human perception and response. Our objectives are threefold: (i) to formulate an adaptive OVM that sustains oscillatory dynamics matching experimental observations, (ii) to fit this model to empirical ring-road data and quantify the improvement over baseline formulations, and (iii) to interpret the fitted parameters in terms of plausible driver behavior.

2 Prior Work

Our work builds upon established microscopic traffic models for phantom jam formation. The foundational Optimal Velocity Model (OVM), introduced by Bando et al. [1], describes drivers adjusting their velocity toward a headway-dependent target speed. While conceptually elegant, the OVM assumes instantaneous driver response, which is physiologically unrealistic.

Bando et al. [4] addressed this limitation by introducing an explicit reaction delay τ. The resulting Delayed OVM demonstrated that finite reaction times significantly reduce traffic stability, amplifying perturbations even in regimes that are stable under instantaneous response. This established delayed feedback as a key mechanism in phantom traffic formation.

Several extensions have since incorporated additional behavioral realism. The Full Velocity Difference Model (FVDM) [5] augments the OVM with a term proportional to the relative velocity between successive vehicles, enabling drivers to anticipate closing or separating gaps. The Generalized Force Model (GFM) [6] reinterprets vehicle acceleration as the result of generalized interaction forces, combining attraction toward a desired velocity with repulsion from unsafe proximity to the leader. Both models improve upon the OVM's stability characteristics but retain constant driver sensitivity.

Experimental validation of the OVM was provided by a controlled ring-road experiment conducted by Sugiyama et al. [3], which demonstrated that traffic jams can arise spontaneously in the absence of bottlenecks. This experiment closely matched the assumptions of the OVM and provided empirical support for microscopic phantom traffic models.

3 Adaptive OVM

We model N identical vehicles indexed by $i = 1, \ldots, N$ on a ring road of length L, imposing periodic boundary conditions

$$x_{N+1}(t) = x_1(t) + L$$

so that headways sum to the ring length,

$$\sum_{i=1}^{N} h_i = L.$$

The state variables are position $x_i(t)$ and velocity $v_i(t)$. The classical OVM [1] describes car-following as a system of coupled ODEs in which drivers continuously adjust their velocity toward a desired speed determined by the gap to the vehicle ahead. We adopt the headway

$$h_i(t) = x_{i+1}(t) - x_i(t),$$

and the relative velocity

$$\Delta v_i(t) = v_{i+1}(t) - v_i(t),$$

which is negative when vehicle i is approaching a slower leader.

The adaptive-sensitivity OVM we analyze is

$$\dot{x}_i(t) = v_i(t),$$
$$\dot{v}_i(t) = a\big(h_i(t), \Delta v_i(t)\big) \left[V\big(h_i(t)\big) - v_i(t)\right].$$

where $a(h_i, \Delta v_i)$ is a state-dependent sensitivity function discussed below, and $V(h)$ is the optimal velocity function. We adopt the standard hyperbolic tangent form [1]

$$V(h) = \frac{V_{\max}}{2} \left[\tanh(h - h_c) + \tanh(h_c)\right],$$

which smoothly interpolates between zero velocity at small headways and a maximum free-flow speed $V_{\max}$ at large headways, with h_c controlling the critical headway at which drivers transition between congested and free-flow regimes.

3.1 Adaptive Sensitivity Function

In the classical OVM, the sensitivity parameter a is constant, implying that drivers react with identical urgency regardless of traffic conditions. Empirical observations suggest otherwise: drivers exhibit heightened responsiveness when gaps are small (small headway) or closing rapidly (large velocity difference with the car in front) [2]. We propose a sensitivity function that captures this behavior:

$$a(h, \Delta v) = a_0 \left[1 + \beta_h \exp\left(-\frac{h - h_c}{\sigma_h}\right) + \beta_v \max(0, -\Delta v)\right],$$

where a_0 is the baseline sensitivity governing relaxed driving conditions, $\beta_h \geq 0$ amplifies sensitivity when headway falls below the critical value h_c, $\sigma_h > 0$ controls the spatial scale over which headway-dependent sensitivity decays, and $\beta_v \geq 0$ increases sensitivity when approaching a slower leader ($\Delta v < 0$). The exponential headway term ensures that sensitivity increases smoothly as gaps shrink, while the velocity term provides asymmetric response since it is known that drivers react more strongly to closing gaps than to opening ones [2]. This asymmetry is consistent with the heightened attention drivers allocate to collision-avoidance scenarios.

3.2 Delayed Adaptive OVM

Human reaction time introduces a fundamental delay between perception and response. Following [4], we incorporate an explicit time delay $\tau > 0$ into the adaptive model. We define the Delayed Adaptive OVM (DA-OVM) as

$$\dot{x}_i(t) = v_i(t),$$

$$\dot{v}_i(t) = a\big(h_i(t-\tau), \Delta v_i(t-\tau)\big) \left[V\big(h_i(t-\tau)\big) - v_i(t)\right].$$

where drivers base their acceleration decisions on traffic conditions observed τ seconds in the past, while responding to their current velocity. This is unlike the FVDM, which adds a velocity difference term directly to the acceleration, the DA-OVM modulates driver responsiveness while preserving the OVM structure and introducing state-dependent gain.This formulation yields a system of delay-differential equations (DDEs) with seven free parameters:

$$\boldsymbol{\theta} = (V_{\max}, h_c, a_0, \beta_h, \sigma_h, \beta_v, \tau).$$

3.3 Parameter Estimation

We estimate the parameter vector $\boldsymbol{\theta}$ by minimizing am objective function that balances trajectory accuracy with oscillation amplitude matching. Given velocity data $\{v_i^{\text{obs}}(t_k)\}$ for $i = 1, \ldots, 20$ vehicles at monotonically increasing times t_k, [3], we define the objective

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{1}{NM} \sum_{i=1}^{N} \sum_{k=1}^{M} \big(v_i(t_k; \boldsymbol{\theta}) - v_i^{\text{obs}}(t_k)\big)^2 + \omega\big(\sigma_{\text{obs}} - \sigma_{\text{pred}}(\boldsymbol{\theta})\big)^2,$$

where $v_i(t_k; \boldsymbol{\theta})$ denotes the simulated velocity, σ_{obs} and σ_{pred} are the standard deviations of observed and predicted velocities respectively (of all 20 cars over a window of 5 seconds), and $\omega > 0$ weights the oscillation penalty.

The first term is the mean squared error (MSE), which measures the accuracy of the trajectory predictions of the model to the data. The second term penalizes the models that underpredict velocity variance, preventing the optimizer from converging to solutions that capture mean behavior while damping out fluctuations. We set $\omega = 0.3$ based on the preliminary sensitivity analysis.

The objective function $\mathcal{L}(\boldsymbol{\theta})$ can become undefined when certain parameter choices cause the simulation to diverge, making gradient-based optimization unsuitable. We instead employ differential evolution [7], a global optimization algorithm well-suited to problems with multiple local minima. The optimization proceeds as follows:

1. We sample candidate parameter vectors uniformly from the bounds specified in Table 1.
2. For each candidate $\boldsymbol{\theta}$, simulate the DA-OVM system from the experimental initial conditions.
3. Compute the objective $\mathcal{L}(\boldsymbol{\theta})$ by comparing simulated velocities to the experimental data.
4. Generate new candidates through mutation and crossover operations, retaining those that reduce the objective value.
5. Repeat steps 2–4 until the objective improves by less than 10^{-5} between iterations, or until 200 iterations have elapsed.
6. Apply a local optimizer (L-BFGS-B) to finetune the best solution.

We use the differential_evolution package from scipy.optimize to implement the algorithm. All models were fitted using the same procedure to ensure fair comparison. Candidate evaluations were parallelized across all available CPU cores, with total fitting time for the seven-parameter DA-OVM of approximately 20 minutes on a 28-core Intel Xeon Gold processor with 400GB of RAM (although it was not close to limited by memory availability). The parameter bounds used in this study are summarized in Table 1.

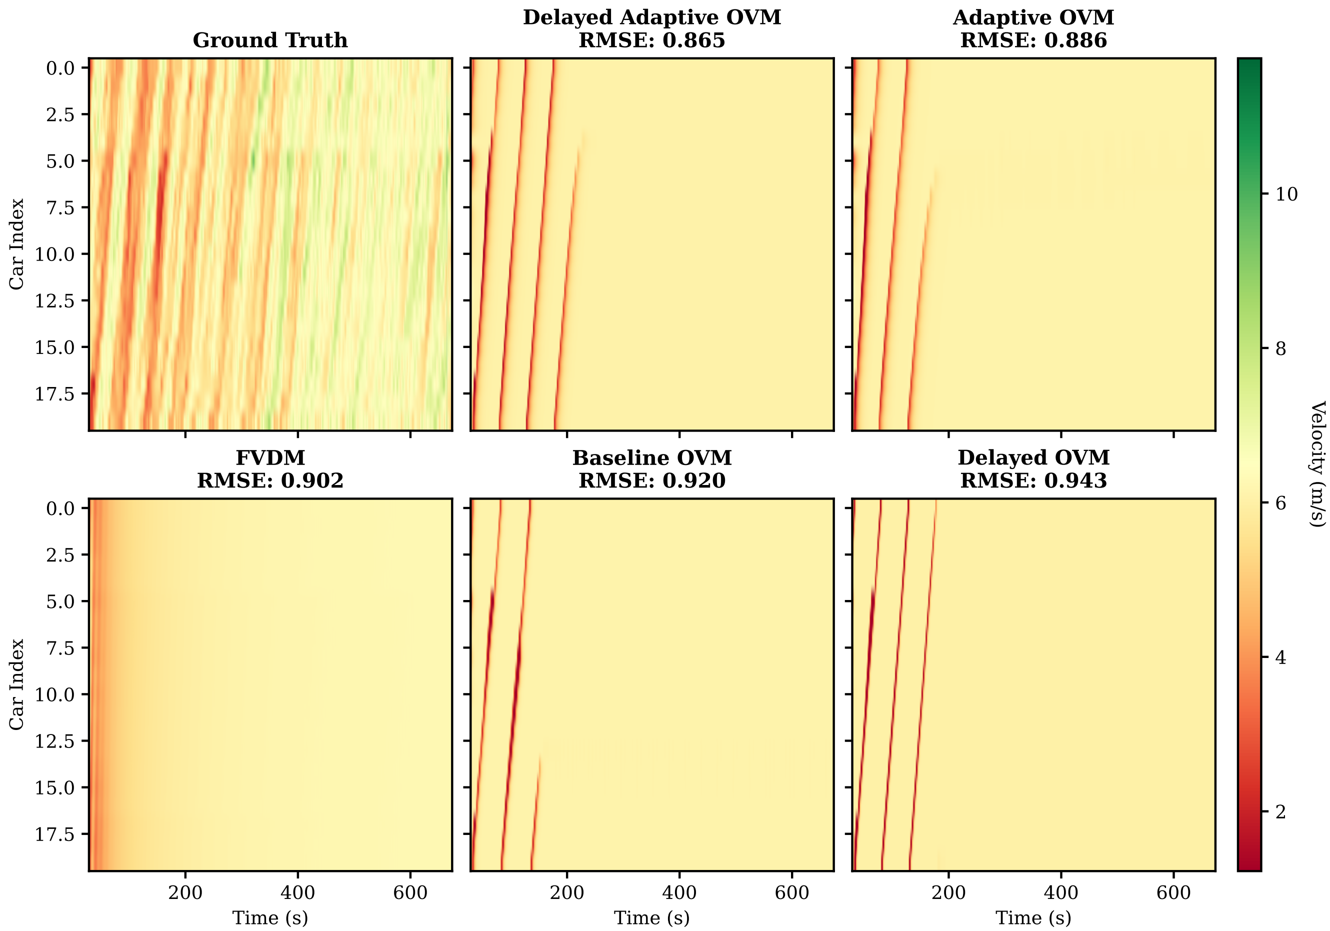

Figure 1: Comparison of the car-velocity-time graphs of 5 fitted models and the ground truth. Our models perform better in terms of RMSE and the reconstruction of the periodic dynamics of the cars.

Param	Low	High	Interpretation
$V_{\max}$	5	15	Maximum velocity (m/s)
h_c	3	15	Critical headway (m)
a_0	0.1	2.0	Baseline sensitivity (s^{-1})
β_h	0	3	Headway sensitivity
σ_h	1	10	Headway decay (m)
β_v	0	1	Velocity sensitivity (s/m)
τ	0.1	2.0	Reaction delay (s)

Table 1: Parameter bounds for DA-OVM optimization.

3.4 Stability and Damping

For the classical Optimal Velocity Model (OVM), small perturbations in vehicle spacing or speed decay over time, causing the system to relax toward homogeneous motion. As a result, the OVM overdamps disturbances and fails to sustain stop–and–go waves beyond a short transient [1].

Introducing an explicit reaction delay reduces this damping by causing drivers to respond to outdated traffic conditions [4]. Delayed OVM variants therefore exhibit longer-lived oscillations, but delay alone is insufficient to maintain oscillations at experimentally observed amplitudes.

The Delayed Adaptive OVM (DA-OVM) preserves linear stability while modifying the damping behavior in a state-dependent manner. When vehicles approach closely or close in on a slower leader, adaptive sensitivity increases driver responsiveness and counteracts the overdamping present in constant-sensitivity models. This nonlinear feedback allows oscillations to persist over multiple cycles, reproducing the sustained periodic structure of phantom traffic observed in ring-road experiments.

3.5 Numerical Integration

The system of coupled ODEs is integrated numerically using the explicit Runge-Kutta method of order 4(5) with adaptive step size control, implemented via `solve_ivp` from SciPy. Initial positions $x_i(t_0)$ and velocities $v_i(t_0)$ are extracted directly from the experimental data. The state vector $\mathbf{y} = (x_1, \ldots, x_N, v_1, \ldots, v_N)^T$ has dimension $2N = 40$.

At each integration step, headways are computed with periodic boundary conditions:

$$h_i = (x_{i+1} - x_i) \mod L - \ell,$$

where $\ell = 3.885$ m is the vehicle length and $L = 2\pi \times 50$ m is the circuit length. The modular arithmetic ensures correct headway computation when vehicle indices wrap around the ring.

For models without delay (Baseline OVM, Adaptive OVM, FVDM, GFM), we use the RK45 integrator directly. For models with delay (Delayed OVM, DA-OVM), the system becomes a set of delay-differential equations. These are integrated using the method of steps with a fixed time step $\Delta t = 0.1$ s, maintaining a history buffer of past states to evaluate the delayed terms $h_i(t-\tau)$ and $\Delta v_i(t-\tau)$. Euler stepping is used within this framework for simplicity, as `solve_ivp` does not natively support delayed arguments.

4 Results

We evaluate the Delayed Adaptive OVM (DA-OVM) against established car-following models using trajectory data from the Sugiyama ring-road experiment. [3]. All models are fitted using the same differential evolution procedure described in section 3.3, ensuring fair comparison.

4.1 Model Comparison

Table 2 summarizes the performance of six car-following models. The DA-OVM achieves the lowest RMSE (0.865 m/s) and highest R^2 (0.181), representing a 6% improvement in RMSE over the baseline OVM and a 4% improvement over the next-best model (Adaptive OVM, also ours).

Among the baseline models (from prior work), the FVDM performs best, suggesting that explicit dependence on relative velocity Δv is important for capturing traffic dynamics. However, all constant-sensitivity models exhibit Std Ratios below 0.62, indicating that they underpredict velocity fluctuations by at least 38%. The velocity fluctuates in an oscillatory manner, which isn't reflected well by the baseline OVM models as they dampen oscillations too aggressively which stabilizes the traffic. The Delayed OVM has the longest oscillatory motion amongst the baseline OVMs (see Figure 2 and Figure 1).

The Adaptive OVM improves upon the baseline by allowing state-dependent sensitivity, achieving an R^2 of 0.140 compared to 0.074 for the baseline. It models the oscillatory motion of the empirical data (very noticeable in Figure 2) well for 2 cycles before stabilizing.

By combining the longer oscillatory motion of the Delayed OVM with the best fitted Adaptive OVM, we created the DA-OVM. We found that by adding reaction delay to the adaptive OVM yielded further gains: the Std Ratio increases from 0.555 to 0.633, indicating that the delay mechanism helps sustain oscillations closer to those observed experimentally. The R^2 value further increases to 0.181, and the mean average error (MAE) is the lowest amongst all the models, indicating a better fit to the empirical data. Figures 1 and 2 show that the DA-OVM captures more than one cycle of phantom traffic more than the traditional model.

4.2 Optimized Parameters

Table 3 reports the fitted parameters for the DA-OVM. The baseline sensitivity $a_0 = 0.185$ s^{-1} corresponds to a relaxation time of approximately 5.4 seconds under normal driving conditions, indicating relatively sluggish response in uncongested flow. The headway amplification $\beta_h = 1.77$ indicates that sensitivity nearly triples when the gap falls below the critical headway $h_c = 5.15$ m. The short decay length $\sigma_h = 1.45$ m implies that this heightened sensitivity is sharply localized to small gaps. The velocity sensitivity $\beta_v = 0.80$ s/m implies that closing on a leader at 2 m/s increases sensitivity by approximately 160%.

The fitted delay $\tau = 0.107$ s is notably shorter than typical human reaction times (0.5–1.5 s) reported in the literature [8]. This suggests that, in the controlled ring-road environment with attentive drivers, perceptual delays are minimal. Alternatively, the short delay may indicate that other model components (particularly the strong velocity sensi-

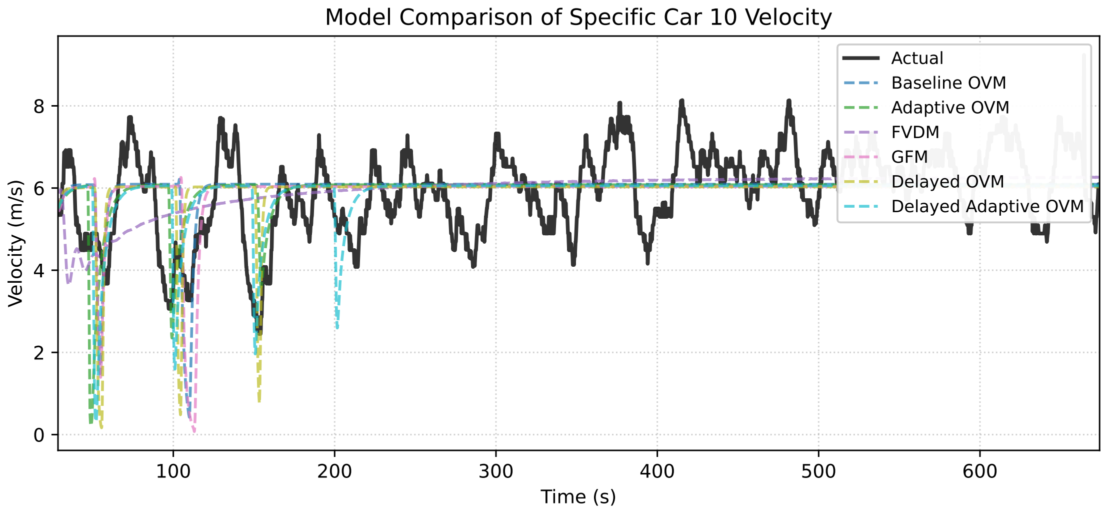

Figure 2: Velocity as a function time of car 10 for the 6 fitted models and the ground truth. The DA-OVM (light blue) sustains the periodic motion for 1 additional oscillation compared to the other models.

Model	**RMSE ↓**	**MAE ↓**	R^2 ↑	**Actual Std**	**Pred Std**	**Std Ratio**
FVDM [5]	0.901767	0.702697	0.109953	0.952282	0.503198	0.528413
Baseline OVM [1]	0.920033	0.716466	0.073531	0.952282	0.493478	0.518206
GFM [6]	0.930069	0.729745	0.053208	0.952282	0.526610	0.552998
Delayed OVM [4]	0.942960	0.739389	0.026781	0.952282	0.581806	0.610960
Adaptive OVM	0.886331	0.696331	0.140163	0.952282	0.528489	0.554971
Delayed Adaptive OVM	**0.865264**	**0.685188**	**0.180553**	0.952282	0.602434	**0.632622**

Table 2: Performance Comparison of Car-Following Models

Parameter	Fitted Value	Unit
$V_{\max}$	6.05	m/s
h_c	5.15	m
a_0	0.185	s^{-1}
β_h	1.77	–
σ_h	1.45	m
β_v	0.80	s/m
τ	0.107	s

Table 3: Fitted DA-OVM parameters.

tivity β_v) already capture much of the destabilizing effect typically attributed to reaction time.

The maximum velocity $V_{\max} = 6.05$ m/s (21.8 km/h) reflects the low-speed regime of the ring-road experiment, while the critical headway $h_c = 5.15$ m corresponds to approximately 1.3 car lengths.

5 Discussion

The fitted parameters reveal interesting insights about driver behavior in congested ring-road traffic. The high value of $\beta_v = 0.80$ s/m suggests that drivers in the experiment were highly attentive to closing speeds. When approaching a leader at $\Delta v = -2$ m/s, sensitivity increases by a factor of $1 + 0.80 \times 2 = 2.6$, dramatically shortening reaction time.

The combination of $\beta_h = 1.77$ and $\sigma_h = 1.45$ m indicates that headway-dependent sensitivity activates sharply only at small gaps. At $h = h_c = 5.15$ m, sensitivity is amplified by factor $1 + 1.77 = 2.77$. At $h = h_c + 3$ m, this amplification drops to $1 + 1.77 \cdot e^{-3/1.45} \approx 1.22$.

The short fitted delay $\tau = 0.107$ s suggests that the adaptive sensitivity function already captures much of the destabilizing dynamics typically attributed to reaction time. The strong β_v term, in particular, creates implicit delay-like effects by amplifying responses to velocity differences.

The present study assumes identical drivers and a single-lane ring-road geometry, neglecting heterogeneity, lane changes, and stochastic effects present in real traffic. The proposed sensitivity function is phenomenological rather than derived from first principles, and alternative functional forms may yield comparable performance. Additionally, the model does not enforce explicit physical braking limits, which may become relevant under extreme conditions. These limitations suggest caution when extrapolating the results beyond controlled experimental settings.

As traffic systems move toward mixed human–autonomous operation, accurate car-following models play a critical role in ensuring both safety and stability. Autonomous vehicles that fail to account for adaptive human response or delayed feedback may inadvertently amplify traffic instabilities. The proposed framework offers a testbed for studying stability-aware control strategies that suppress phantom traffic while remaining compatible with human driving behavior.

6 Conclusion

We have introduced the Delayed Adaptive Optimal Velocity Model (DA-OVM), which augments the classical OVM by allowing driver sensitivity to depend on headway and relative velocity while explicitly accounting for reaction delay. When fitted to empirical ring-road data, the DA-OVM consistently outperformed classical and extended OVM variants, achieving lower prediction error and sustaining oscillatory dynamics closer to those observed experimentally. Unlike constant-sensitivity models, the DA-OVM reproduces the persistent, periodic structure of phantom traffic jams rather than overdamping velocity fluctuations.

Analysis of the fitted parameters suggests that human drivers exhibit strongly adaptive responsiveness, reacting significantly more aggressively when gaps are small or closing rapidly. The short fitted reaction delay further indicates that adaptive sensitivity may capture destabilizing effects traditionally attributed to delayed response.

Future work will focus on analytical stability analysis of the DA-OVM, extension to heterogeneous driver populations, and investigation of how adaptive sensitivity interacts with automated or partially automated vehicles in mixed traffic.

References

[1] Masaki Bando, Kazumasa Hasebe, Akihiro Nakayama, Akira Shibata, and Yuki Sugiyama. Dynamical model of traffic congestion and numerical simulation. *Physical Review E*, 51(2):1035–1042, 1995.

[2] Martin Treiber and Dirk Helbing. Memory effects in microscopic traffic models and wide scattering in flow–density data. *Physical Review E*, 68(4):046119, 2003.

[3] Yuki Sugiyama, Minoru Fukui, Macoto Kikuchi, Kazumasa Hasebe, Akihiro Nakayama, Katsuhiro Nishinari, Shin-ichi Tadaki, and Satoshi Yukawa. Traffic jams without bottlenecks—experimental evidence for the physical mechanism of the formation of a jam. *New Journal of Physics*, 10(3):033001, 2008.

[4] Masaki Bando, Kazumasa Hasebe, Akihiro Nakayama, Akira Shibata, and Yuki Sugiyama. Analysis of optimal velocity model with explicit delay. *Physical Review E*, 58(5):5429–5435, 1998.

[5] Rui Jiang, Qing-Song Wu, and Zhen Zhu. Full velocity difference model for a car-following theory. *Physical Review E*, 64(1):017101, 2001.

[6] Dirk Helbing and Benno Tilch. Generalized force model of traffic dynamics. *Physical Review E*, 58(1):133–138, 1998.

[7] Rainer Storn and Kenneth Price. Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces. *Journal of Global Optimization*, 11(4):341–359, 1997.

[8] Michael Green. "How long does it take to stop?" methodological analysis of driver perception–brake times. *Human Factors*, 42(2):195–216, 2000.