
RACER: Rational Artificial Intelligence Car-following-model Enhanced by Reality

Tianyi Li*

Department of Civil, Environmental, and Geo- Engineering
University of Minnesota
Minneapolis, MN 55414
li001915@umn.edu

Raphael E. Stern

Department of Civil, Environmental, and Geo- Engineering
University of Minnesota
Minneapolis, MN 55414
rstern@umn.edu

Abstract

This paper introduces RACER, the Rational Artificial Intelligence Car-following model Enhanced by Reality, a cutting-edge deep learning car-following model, which satisfies partial derivative constraints that are necessary to maintain physical feasibility, designed to predict Adaptive Cruise Control (ACC) driving behavior. Unlike conventional car-following models, RACER effectively integrates Rational Driving Constraints (RDC), crucial tenets of actual driving, resulting in strikingly accurate and realistic predictions. Notably, its adherence to the RDC, registering zero violations, is in stark contrast to other models. This study incorporates physical constraints within AI models, especially for obeying rational behaviors in transportation. The versatility of the proposed model, including its potential to incorporate additional derivative constraints and broader architectural applications, enhances its appeal and broadens its impact within the scientific community.

1 Introduction

Advancements in vehicle automation are reshaping the transportation landscape, with a dual impact on traffic flow and stability. While some research underscores the benefits of automated vehicle (AV) technologies like enhanced traffic flow (Tan et al., 1998; Wang et al., 2022) and speed harmonization (Learn et al., 2017), other studies point to potential drawbacks. Notably, adaptive cruise control (ACC) vehicles, which are the first generation of AV, could reduce highway throughput (Shang and Stern, 2021). This underscores the need for accurate car-following models that capture the nuanced behavior of ACC and AV technologies (Talebpour and Mahmassani, 2015; Shang and Stern, 2021).

Various modeling approaches have emerged to understand vehicle-level dynamics in automated systems (Talebpour and Mahmassani, 2016; Milanés and Shladover, 2014). Most adapt traditional car-following models for new automated features. As lower-level automation features like ACC become widespread, understanding their impact on traffic dynamics becomes crucial (Gunter et al., 2020; Shang and Stern, 2021).

*T. Li acknowledges the support of the Dwight David Eisenhower Graduate Fellowship from the Federal Highway Administration

Car-following models have seen significant evolution, incorporating technologies like deep learning (Wang et al., 2017; Mo et al., 2021). These models generally fall into physics-based, data-driven, and physics-guided AI categories (Mo et al., 2021). Each has its limitations, such as oversimplification or interpretability issues (Wang et al., 2017; Raissi et al., 2019).

Recently, hybrid models combining physics and data-driven approaches have emerged (Raissi et al., 2019; Mo et al., 2021). These promising frameworks integrate domain-specific insights with machine learning capabilities. However, they often overlook Rational Driving Constraints (RDC) (Wilson and Ward, 2011), critical for understanding realistic driving behavior.

Our work addresses these gaps by integrating RDCs and physical constraints into a neural network-based car-following model. The aim is to produce a more reliable and interpretable model that advances the capabilities of existing approaches, setting a foundation for safer autonomous driving systems. Our contributions are threefold: 1) a novel methodology that embeds RDCs; 2) superior performance compared to existing models; and 3) an analysis showing compliance with RDC constraints, a first in machine learning-based car-following models.

2 Modeling Car-following Behavior

This section introduces our neural network-based car-following model, specifically designed to comply with RDCs and to adapt to various driving conditions. We detail the model architecture, its training process, and its alignment with the RDC.

A car-following model (CFM) is a function f_θ , parameterized by θ , that maps state variables o to actions a (longitudinal accelerations):

$$f_\theta : o \rightarrow a \quad (1)$$

Our model primarily focuses on three variables: spacing $s(t)$, relative speed (the difference between lead vehicle speed and following vehicle speed) $\Delta v(t)$, and the subject vehicle’s velocity $v(t)$, constituting the state vector $o = (s(t), \Delta v(t), v(t))$. The acceleration $a(t)$ is represented as a second-order ordinary differential equation:

$$\ddot{x}(t) = f_\theta(s(t), \Delta v(t), v(t)) \quad (2)$$

Here, x denotes position on the road, $v(t) = \dot{x}(t)$, $\Delta v(t) = \dot{s}(t)$, and $a(t) = \ddot{x}(t)$.

2.1 Deep Learning-Based Car-Following Model with Rational Driving Constraints

Previous machine learning-based car-following models excel in certain metrics like spacing RMSE but often overlook the integration of RDCs (Panwai and Dia, 2007; Huang et al., 2018; Wang et al., 2017; Mo et al., 2021; Naing et al., 2022; Ma et al., 2023; Zhu et al., 2018). We propose “RACER”, a model that integrates RDCs to guide the predictions toward safety and realism.

RDCs serve as foundational constraints that embody the behavioral laws to which rational drivers conform (Wilson and Ward, 2011; Stern et al., 2018). For example, decelerate when the speed is high. This condition is mathematically represented as a non-positive derivative of acceleration with respect to speed. These are mathematically expressed as:

$$\frac{da}{dv} \leq 0, \quad \frac{da}{ds} \geq 0, \quad \frac{da}{d(\Delta v)} \geq 0 \quad (3)$$

Implementing these constraints in learning-based car-following models is essential as they warrant that the model’s predictions align with basic safe and rational driving principles. Nonetheless, enforcing these constraints, especially in intricate models like deep neural networks, is a challenging task. The proposed solution in the provided code introduces a novel approach to enforce RDCs by incorporating them into the loss function of the model. The details of the algorithm are presented

in Algorithm 1. Here, NN_{seq} represents the sequence handling neural network, $\text{process}(\cdot)$ signifies additional processing steps (that could vary depending on the exact network architecture used), and $\text{combine}(\cdot)$ refers to the operation that combines the processed sequence and physical inputs.

Algorithm 1 Enforcing Rational Driving Constraints

- 1: **Input:** Sequence of vehicle states X_{seq} , Physical vehicle states X_{phy}
 - 2: **Output:** Predicted acceleration a_{pred} , Loss \mathcal{L}
 - 3: $Z_{\text{seq}} \leftarrow NN_{\text{seq}}(X_{\text{seq}})$
 - 4: $Z'_{\text{seq}} \leftarrow \text{process}(Z_{\text{seq}})$
 - 5: $Z'_{\text{phy}} \leftarrow \text{process}(X_{\text{phy}})$
 - 6: $a_{\text{pred}} \leftarrow \text{combine}(Z'_{\text{seq}}, Z'_{\text{phy}})$
 - 7: Calculate gradients $\frac{\partial a_{\text{pred}}}{\partial v}$, $\frac{\partial a_{\text{pred}}}{\partial s}$, $\frac{\partial a_{\text{pred}}}{\partial r}$
 - 8: $RDC_{\text{speed}} \leftarrow \text{ReLU}\left(\frac{\partial a_{\text{pred}}}{\partial v}\right)$
 - 9: $RDC_{\text{spacing}} \leftarrow \text{ReLU}\left(-\frac{\partial a_{\text{pred}}}{\partial s}\right)$
 - 10: $RDC_{\text{relative speed}} \leftarrow \text{ReLU}\left(-\frac{\partial a_{\text{pred}}}{\partial r}\right)$
 - 11: $\mathcal{L} \leftarrow \text{MSE}(a_{\text{pred}}, a_{\text{true}}) + \lambda \cdot (RDC_{\text{speed}} + RDC_{\text{spacing}} + RDC_{\text{relative speed}})$
 - 12: **return** a_{pred} , \mathcal{L}
-

3 Numerical Experiments

3.1 Data Description

Our analysis employs a principal dataset derived from a sequence of car-following experiments conducted by Gunter et al. (Gunter et al., 2020). This dataset is amassed using a variety of commercially available vehicles equipped with ACC systems. Each of these ACC-activated vehicles adheres to a uniform testing procedure, where a leading vehicle traverses at a pre-established speed sequence for a set duration at each pace. The ACC vehicle, while trailing the lead vehicle, has its ACC active throughout the duration of the experiment.

3.2 Analysis of the Numerical Experiments

We evaluate four distinct models: Optimal Velocity Relative Velocity (OVRV), Neural Network (NN), Physics Informed Neural Network (PINN), and Rational Neural Networks (Rational NN). As recommended by Punzo and Montanino (Punzo and Montanino, 2016), we prefer to compare model performance using the cumulative inter-vehicle spacing rather than the instantaneous values. Consequently, we take into account the cumulative error for the temporal evolution of states. Using our models’ acceleration predictions, we reconstruct position and velocity trajectories, thus the spacing plot in Figure 1 uses the proposed model as a controller. These trajectories (spacing and speed profiles) for the following vehicles are derived based on kinematic dynamics, using a time step of $\Delta t = 0.1$ s—aligned with the experimentally gathered data:

$$\begin{bmatrix} s \\ v \end{bmatrix}_{t+\Delta t} = \begin{bmatrix} s \\ v \end{bmatrix}_t + \begin{bmatrix} v_l - v \\ a_{\text{pred}} \end{bmatrix}_t \Delta t, \quad (4)$$

Our evaluation indicates that the Rational NN model consistently outperforms all other models, notably excelling in capturing the dynamic behavior of acceleration and deceleration phases. Figures 1a and 1c substantiate this superiority, with the Rational NN model providing the closest adherence to actual vehicular behavior. Importantly, the model avoids the overshooting problem commonly observed in other models, offering a more reliable and safe predictive capability. Our model’s predictions were assessed for compliance with RDCs across three parameters: velocity (v), spacing (s), and relative speed (r). Unlike other models that failed in RDC compliance, ours recorded zero

Table 1: Root Mean Squared Errors for Different Models.

	RACER	OVRV	NN	PINN
Acceleration (m/s²)	0.099	0.111	0.115	0.111
Speed (m/s)	0.152	0.173	0.237	0.322
Spacing (m)	0.298	1.485	0.559	0.415




Figure 1: Model performance in terms of acceleration, spacing, and speed.

violations, as shown in Figure 2. This underscores the model’s effectiveness and the advantages of incorporating domain-specific knowledge for realistic and reliable predictions.

4 Discussion & Conclusion

Our experiments highlight that physically constrained neural networks outperform pure neural network models in handling nonlinear problems and providing more accurate predictions than conventional models like OVRV. While integrating a physical model like OVRV as a loss function can sometimes worsen performance, the careful inclusion of RDCs in our model yielded enhanced accuracy. Not only did our model surpass others in predictive power, but it also generated more rational and safer driving trajectories, thus enhancing overall vehicle safety.

We have not yet evaluated our model in the context of human driving data. Human driving behavior is undoubtedly more complex, diverse, and occasionally may not strictly adhere to the RDC. However, our proposed model, which effectively combines physical-guided AI with car-following principles, exhibits enormous potential to inform and guide human drivers, thus making driving safer and more rational. The model also holds promise for testing across a wide range of ACC vehicle driving scenarios as well as human driving situations. Future research could explore this exciting avenue and contribute to the progressive journey toward achieving safer, smarter, and more efficient transportation systems.




Figure 2: Visual Comparisons of Speed Predictions and Rule Violations for Three Models: a) LSTM Model, b) PINN Model, and c) RACER Model. The green and red dots denote predictions conforming to and violating the established rules, respectively.

References

- Gunter, G., Gloudemans, D., Stern, R. E., McQuade, S., Bhadani, R., Bunting, M., Delle M., M. L., Lysecky, R., Seibold, B., Sprinkle, J., et al. (2020). Are commercially implemented adaptive cruise control systems string stable? *IEEE Transactions on Intelligent Transportation Systems*.
- Huang, X., Sun, J., and Sun, J. (2018). A car-following model considering asymmetric driving behavior based on long short-term memory neural networks. *Transportation research part C: emerging technologies*, 95:346–362.
- Learn, S., Ma, J., Raboy, K., Zhou, F., and Guo, Y. (2017). Freeway speed harmonisation experiment using connected and automated vehicles. *IET Intelligent Transport Systems*, 12(5):319–326.
- Ma, L., Qu, S., Song, L., Zhang, Z., and Ren, J. (2023). A physics-informed generative car-following model for connected autonomous vehicles. *Entropy*, 25(7):1050.
- Milanés, V. and Shladover, S. E. (2014). Modeling cooperative and autonomous adaptive cruise control dynamic responses using experimental data. *Transportation Research Part C: Emerging Technologies*, 48:285–300.
- Mo, Z., Shi, R., and Di, X. (2021). A physics-informed deep learning paradigm for car-following models. *Transportation research part C: emerging technologies*, 130:103240.
- Naing, H., Cai, W., Nan, H., Tiantian, W., and Liang, Y. (2022). Dynamic data-driven microscopic traffic simulation using jointly trained physics-guided long short-term memory. *ACM Transactions on Modeling and Computer Simulation*, 32(4):1–27.
- Panwai, S. and Dia, H. (2007). Neural agent car-following models. *IEEE Transactions on Intelligent Transportation Systems*, 8(1):60–70.
- Punzo, V. and Montanino, M. (2016). Speed or spacing? cumulative variables, and convolution of model errors and time in traffic flow models validation and calibration. *Transportation Research Part B: Methodological*, 91:21–33.
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707.
- Shang, M. and Stern, R. E. (2021). Impacts of commercially available adaptive cruise control vehicles on highway stability and throughput. *Transportation Research Part C: Emerging Technologies*, 122:102897.
- Stern, R. E., Cui, S., Monache, M. L. D., Bhadani, R., Bunting, M., Churchill, M., Hamilton, N., Haulcy, R., Pohlmann, H., Wu, F., Piccoli, B., Seibold, B., Sprinkle, J., and Work, D. B. (2018). Dissipation of stop-and-go waves via control of autonomous vehicles: Field experiments. *Transportation Research Part C: Emerging Technologies*, 89:205 – 221.
- Talebpour, A. and Mahmassani, H. S. (2015). Influence of autonomous and connected vehicles on stability of traffic flow. Technical report.
- Talebpour, A. and Mahmassani, H. S. (2016). Influence of connected and autonomous vehicles on traffic flow stability and throughput. *Transportation Research Part C: Emerging Technologies*, 71:143–163.
- Tan, H.-S., Rajamani, R., and Zhang, W.-B. (1998). Demonstration of an automated highway platoon system. In *Proceedings of the 1998 American control conference. ACC (IEEE Cat. No. 98CH36207)*, volume 3, pages 1823–1827. IEEE.
- Wang, S., Stern, R., and Levin, M. (2022). Optimal control of autonomous vehicles for traffic smoothing. *IEEE Transactions on Intelligent Transportation Systems*, 23(4):3842–3852.
- Wang, X., Jiang, R., Li, L., Lin, Y., Zheng, X., and Wang, F.-Y. (2017). Capturing car-following behaviors by deep learning. *IEEE Transactions on Intelligent Transportation Systems*, 19(3):910–920.

Wilson, R. E. and Ward, J. A. (2011). Car-following models: fifty years of linear stability analysis—a mathematical perspective. *Transportation Planning and Technology*, 34(1):3–18.

Zhu, M., Wang, X., and Wang, Y. (2018). Human-like autonomous car-following model with deep reinforcement learning. *Transportation research part C: emerging technologies*, 97:348–368.