



# Teaching Neural Networks the Rules on the Road

## How Physics-Informed Neural Networks Could Change How We Manage Road Traffic Systems

*AiMTT tutorial #4*

*Published: July 2026*

### **The authors**

*Ying-Chuan Ni is a researcher at the Traffic Engineering Group at ETH Zurich. He is supervised by Dr. Anastasios Kouvelas, and Dr. Michail Makridis.*

*Theivaprakasham Hari is a researcher at TU Delft's DAIMoND Lab. His supervisors are Dr. Ir. Winnie Daamen, Dr. Yanan Xin, and Prof. Dr. Ir. Serge Hoogendoorn.*

## **Introduction**

In 1955, the British mathematician James Lighthill and his colleague Gerald Whitham published a paper describing how vehicles on a long, crowded road behave like a flowing fluid.<sup>1</sup> A year later, the American engineer Paul Richards reached the same conclusion independently.<sup>2</sup> Their combined work, now known as the Lighthill-Whitham-Richards (LWR) model, is still taught in every transport planning or traffic engineering course in the world. It explains, with a single equation, why a small slowdown at the front of a queue can travel backward as a wave and reach drivers who are kilometers away from the original problem. Throughout most of the twentieth century, traffic engineers and researchers

---

<sup>1</sup> Michael James Lighthill and Gerald Beresford Whitham. *On kinematic waves II. A theory of traffic flow on long crowded roads*. Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences, 229(1178):317–345, 5 1955.

<sup>2</sup> Paul I. Richards. *Shock Waves on the Highway*. Operations Research, 4(1):42–51, 2 1956.

relied heavily on this model—or what we call in this article the physics approach—to study congestion propagation phenomena.

Although this physics-based perspective is elegant and widely applicable, the problem is that roads in reality are messier than equations. Drivers get distracted, weather changes, a delivery van double parks, school zones flood with activity in the afternoon. To account for these unpredictable

factors, urban planners turned to the second way of thinking based on empirical data. By leveraging networks of sensors, traffic cameras, and mobile phone tracking, computers can monitor real-time conditions and identify underlying patterns. This is the world of data-driven machine learning approaches.

Both methodologies have different vulnerabilities. Pure physics models can often be oversimplified to match the real world. Conversely, pure data models can be very accurate on the days they were trained on but sometimes produce physically-impossible outcomes when conditions change because they have no knowledge that cars cannot spontaneously appear out of nowhere, nor do they disappear mysteriously, as illustrated in Figure 1. In addition, a road has a maximum capacity for vehicles to pass through and a maximum number of vehicles that it can accommodate at a time. All these are physics laws of the road. What if we could give a neural network a copy of the rule book before we let it loose on the data?

This article explains a novel class of machine learning approaches, the physics-informed neural network, and explores how it can be applied to solve traffic-related problems. We also highlight several of its challenges that require careful consideration.




Figure 1: Comics illustrating the violation of conservation law in traffic flow theory.

## What is a Physics-Informed Neural Network?

A neural network (NN) is a piece of software loosely inspired by the way neurons in the brain connect and communicate. By exposing an NN to many examples, it gradually adjusts its internal parameters until it can make accurate predictions on previously unseen examples. The clever twist in a physics-informed neural network (PINN), first introduced in 2017 by Maziar Raissi, Paris Perdikaris and George Karniadakis, is to embed the laws of physics directly into the network's training process.<sup>3</sup> Whenever the network makes a prediction, it is checked twice. Once against the real measurements, and once against the equation that the system is supposed to obey. If the prediction violates the equation, the network is nudged back into line.

To fully appreciate this innovation, we must look at the fundamental flaws of pure NNs. Standard machine learning (ML) models are exceptional at memorizing patterns within the data they have already seen. However, they struggle severely when asked to extrapolate into completely new situations. A pure data-driven algorithm might perfectly forecast typical weekday commutes but fail completely during a sudden extreme weather event or an unprecedented road closure. Because they rely exclusively on historical trends, standard networks essentially guess blindly when pushed outside their comfort zone.

In contrast, a PINN inherently understands the boundaries of reality. Even if a PINN has never encountered a specific traffic anomaly, it is still aware of the foundational laws of motion. It knows that cars cannot teleport and that traffic shockwaves must travel at mathematically defined speeds. This physical awareness gives PINNs a massive advantage in generalization and reliability over pure NNs. By anchoring the ML process to established scientific laws, developers can build models that remain highly robust in unpredictable real-world environments.

For traffic experts, the practical appeal here is a drastically reduced reliance on massive datasets because the embedded physics constraints naturally fill the knowledge gaps. A few scattered motorway sensors on a motorway, combined with a solid physical traffic model, can be enough for a PINN to reconstruct what is happening between the sensors, including in places where there are no measurements at all. This unique combination of data adaptability and physical

---

<sup>3</sup> M. Raissi, P. Perdikaris, and G.E. Karniadakis. *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, 2 2019.

strictness creates a perfect bridge to our next topic, which explores exactly how we apply these tools to actual road networks.

## **Applications of PINNs in Traffic Flow**

Historically, traffic flow has been considered a matter of physics because it involves the movement and interaction of objects over space and time. Cars on a road behave similarly to particles in gas or fluid. This analogy has led scientists and engineers to use principles from physics to understand how vehicles interact under different conditions, study collective traffic phenomena, model traffic flow mathematically, predict congestion, and design safer and more efficient roads.

Therefore, it is unsurprising that PINN has already been applied to a wide range of problems in the field of traffic flow modeling and control since its first introduction. Now let us look at how it can help us estimate or even predict traffic at two different levels-of-aggregation.

### **PINNs for Macroscopic Traffic States**

A macroscopic traffic state can typically be described by three variables, flow, density, and speed. Due to technological limitations and financial concerns, traffic data collected by state-of-the-art sensors are often sparse and noisy. Traffic state estimation (TSE) is a process that infers the spatiotemporal traffic states on road segments using partially-observed data. It is an essential component in traffic control strategies that operate based on precise traffic state information.

Classical TSE approaches can generally be categorized into two groups: model-driven data assimilation techniques, including Kalman filtering and particle filtering, which are based on physical traffic flow models; and data-driven statistical or ML approaches, such as linear regression, Autoregressive Integrated Moving Average, and Bayesian network. The former relies on a priori knowledge of traffic flow dynamics, requires relatively little input data, and ensures highly interpretable outcomes. However, their performance deteriorates when the underlying models are poorly calibrated or fail to accurately represent traffic phenomena. In contrast, the latter mainly exploits patterns in historical data rather than relying on theoretical assumptions. Their major drawbacks are the high computational cost or the need for abundant data

associated with training and their sensitivity to irregular incidents or anomalies in the dataset.<sup>4</sup>

PINN appears to be a promising way to combine the strengths of both types of TSE approaches. In this article, we elaborate a PINN use case that utilizes both data and the first-order traffic flow model to estimate densities on a motorway segment.

The density contour plot of a 6-km motorway segment in 30 minutes generated using the Cell-Transmission-Model (CTM) is shown in the left subplot of Figure 2. An accident occurs in the middle of the segment for a period of time, blocking the road and causing a backward propagating jam wave that is eventually dissipated due to low inflow demand.




Figure 2: Density contour plots of a 6-km motorway segment in 30 minutes.

Given that loop detectors are only placed every 450 m, the observed traffic states are those shown in the middle subplot. The PINN task aims to estimate the unknown densities in-between the observed locations, as shown on the right subplot, according to the observed densities altogether. Figure 3 illustrates a deep neural network (DNN) architecture that takes the observed density values as input and outputs the unobserved densities.

---

<sup>4</sup> Toru Seo, Alexandre M. Bayen, Takahiko Kusakabe, and Yasuo Asakura. *Traffic state estimation on highway: A comprehensive survey*. Annual Reviews in Control, 43:128–151, 2017.




Figure 3: DNN architecture for traffic state estimation.

Usually, the loss function for the training of a DNN is governed by the differences between the data and the output value at every collocation point, as expressed by the data loss term  $L_{\text{data}}$  in Equation 1. In PINN, the equation is augmented by another physics loss term  $L_{\text{physics}}$  to capture the violation of the PDE. An example of such a PDE could be the well-known LWR continuity equation, also known as the conservation law of traffic flow, as written in Equation 2. The weight factors  $\omega_d$  and  $\omega_p$  can be adjusted for each scenario according to the data quality and to what extent the model describes the physics.

$$L_{\text{total}} = \omega_d L_{\text{data}} + \omega_p L_{\text{physics}} \quad (1)$$

$$\partial k / \partial t + \partial q / \partial x = 0 \quad (2)$$

Besides the LWR model, second-order traffic flow models that aim to better capture the effect of stop-and-go waves, such as the Aw-Rascle-Zhang model, may also be considered for the physics loss term.

Having multiple traffic state data (density contour plots) of this motorway segment, the DNN can be trained to extract the spatiotemporal dependencies between traffic states and in the meantime respect the principle described by the LWR conservation law. By finding a balance between historical data and the law of physics, the DNN is able to reduce the required data size in the training process, mitigate the overfitting problem, and become more robust to noise in the data.

In addition to TSE, PINN can also be applied to traffic prediction problems with a similar problem formulation and physics constraint while considering recurrent neural networks to learn the pattern in sequential data.<sup>5</sup> The improvement in predictive power owing to the inclusion of established physics contributes significantly to real-time traffic control strategies, particularly in situations of incomplete or erroneous data.

It is worth noting that traffic congestion, by its nature, contains large heterogeneity, as explained by the scattering on the congested branch of a fundamental diagram, and can be caused by numerous possible reasons. Congestion can even emerge within an existing jam wave. Furthermore, a practical challenge for traffic prediction using ML approaches is that congested traffic conditions are relatively infrequent compared to free-flow conditions. Despite its better data efficiency than those purely data-driven NNs, training a reliable PINN that can successfully capture the occurrence and propagation of congestion still requires careful preparation of representative data that include all possible types of congestion events, e.g., lane-drop bottleneck, stop-and-go waves, and incidents, etc.

## **PINNs for Microscopic Vehicle Movements and Trajectories**

In addition to macroscopic traffic states, many studies have investigated the use of PINN for microscopic vehicle movements.<sup>6</sup> By combining a deep learning-based model with a car-following model, such as the Intelligent Driver Model (IDM), the Optimal Velocity Model, and the Wiedemann model, to compute the future accelerations, one can reconstruct or even predict future vehicle trajectories.

Figure 4 shows a PINN training architecture for vehicle longitudinal trajectory prediction that uses IDM as the physics model. IDM estimates the next acceleration  $a_{t+1}$  based on the current speed  $v_t$ , headway  $h_t$ , and relative speed  $\Delta v_t$ . The modeled longitudinal position  $x_{t+1}^p$  is then used to calculate the physics

---

<sup>5</sup> Feng Shao, Hu Shao, Xin Wu, Qixiu Cheng, and William H.K. Lam. *A physics-informed machine learning framework for speed-flow prediction: Integrating an S-shaped traffic stream model with deep learning models*. Transportation Research Part C: Emerging Technologies, 180:105362, 11 2025.

<sup>6</sup> Zhaobin Mo, Rongye Shi, and Xuan Di. *A physics-informed deep learning paradigm for car-following models*. Transportation Research Part C: Emerging Technologies, 130:103240, 9 2021; Maosi Geng, Junyi Li, Yingji Xia, and Xiqun (Michael) Chen. *A physics-informed Transformer model for vehicle trajectory prediction on highways*. Transportation Research Part C: Emerging Technologies, 154:104272, 9 2023.

loss  $L_{\text{physics}}$ . The total loss consists of the physics loss  $L_{\text{physics}}$  and the data loss  $L_{\text{data}}$ , which is obtained by comparing the estimate  $x'_{t+1}$  with the historical data  $x^d_{t+1}$ .




Figure 4: PINN training architecture for vehicle longitudinal trajectory prediction.

Compared to car-following behavior, the modeling of drivers' lane-changing behavior remains largely understudied in the field. When the research problem expands from one-dimensional to two-dimensional, the complexity increases significantly because of the interdependencies between the longitudinal vehicle dynamics and the lateral motions. By integrating empirical behavioral patterns with vehicle dynamics constraints, PINN could outperform other conventional ML approaches in predicting lane-changing maneuvers.<sup>7</sup>

A better prediction of vehicle movements through PINN not only helps the development of microscopic traffic simulation tools. With the increasing number of vehicles equipped with self-driving functionality on the road, gaining a better understanding of other vehicles' decision-making mechanisms becomes crucial for their navigation systems. PINN could be used to establish a safe, efficient, and trust-worthy road environment in the era of autonomous driving.

<sup>7</sup> Zhuo Cao, Zuduo Zheng, Mehmet Yildirimoglu, and Shimul Md. Mazharul Haque. *A physics-informed uncertainty quantification framework for deep learningbased real-time lane-change intention prediction*. Transportation Research Part C: Emerging Technologies, 187:105637, 6 2026.

## The Realities and Limitations of PINNs

Despite the growing enthusiasm for hybrid modeling, professionals must understand the inherent limitations of this technology. Primarily, these networks cannot magically fix fundamentally flawed physical models. If developers supply the algorithm with overly simplistic equations that ignore real road friction, driver reaction delays, or harsh weather impacts, the network will rigidly enforce an inaccurate version of reality. The intelligence of a PINN is strictly bounded by the mathematical truth we provide.

Additionally, optimizing these architectures is an immense challenge. As Equation 1 illustrates, the system continuously struggles to satisfy two different masters. It must match actual road observations while strictly following mathematical boundaries. Finding the perfect balance between the data weight and the physics weight requires extensive trial and error. Overemphasizing the physics makes the predictions overly stiff and blind to unique real-world variations. Underemphasizing the equations turns the system right back into an opaque data-crunching machine. Evaluating the physics loss also requires processing complicated gradient calculations, drastically increasing the computational burden and total training hours compared to standard deep learning methods.

Lastly, while these systems are famous for thriving on limited sensor input, they still absolutely require high-quality data for thorough validation. Microscopic vehicle trajectory studies benefit from several massive open-source datasets. Conversely, training and proving the reliability of a macroscopic traffic estimation PINN demands exhaustive, flawless, high-resolution density maps covering large geographical areas. Procuring such comprehensive ground-truth metrics is exceptionally difficult and currently slows down the widespread adoption of PINNs for municipal traffic grids.

## Looking Ahead

Compared to conventional black-box NNs that learn solely from empirical data, PINNs ensure that the results adhere to the known physics laws through incorporating them into the loss function directly, improving the interpretability and generalizability of the network in different scenarios. They are particularly useful when data are sparse or noisy.

Over the past decades, traffic flow has been widely investigated as a physical system whose dynamics are governed by mathematical equations. Many models have been proposed to describe traffic phenomena at the microscopic and

macroscopic levels. This is exactly why PINNs are perfectly suited for the prediction of macroscopic traffic dynamics and microscopic vehicle movements.

However, one should bear in mind that their performance depends on the accuracy of the underlying equations. During training, the model must balance matching observed data with satisfying physical constraints. Overemphasizing physics can make predictions too rigid, while relying too much on data weakens the physics guidance. In addition, enforcing physical laws requires complex gradient calculations, making PINNs more computationally demanding than conventional NNs.

ML approaches have attracted significant attention from traffic flow researchers for many years. Nevertheless, their real-world application is still in its infancy. To bridge the gap between research and practical road traffic management applications and bring these innovative approaches from scientific papers into practice, PINNs are expected to play a pivotal role.