# Driving Style-Aware Car-Following Modeling and Style's Impact on Traffic Breakdown

***Abstract*—Traffic congestion has has become a global challenge with severe economic and safety consequences. It is caused by infrastructure limits but also by diverse human driving behaviors. Traditional car-following (CF) models capture speed and spacing dynamics but limited attention is paid to representing driver heterogeneity. This study addresses this gap by developing a style-aware CF model that jointly learns driving-style embeddings and trajectory prediction within a Transformer framework. The proposed model consistently outperforms the benchmark CF models in long-duration evaluation. Remarkably, the learned style embeddings reveal four interpretable groups along aggressiveness and stability dimensions.**

**The paper also relates the driving styles to macroscopic properties of the traffic stream through platoon-scale simulations. Results show a clear stability–efficiency trade-off: stable–conservative styles buffer congestion but increase fuel consumption through prolonged low-speed operation, while aggressive styles appear fuel-efficient in isolation. Unstable driving style destabilizes the traffic flow the most. Scattered trucks throughout the traffic stream improve stability, whereas clustering trucks increases the probability of a traffic breakdown. These findings highlight how integrating driving style can improve both car-following prediction and its impact on traffic breakdown.**

***Index Terms*—Driving style, Car-following, Traffic congestion, Human factors, Data-driven**

## I. Introduction

Traffic congestion remains a global challenge with serious economic and safety consequences. In the Netherlands, congestion-related economic losses rose by 49% in 2023 compared to 2022 (Rijkswaterstaat, 2024), reflecting the growing severity of traffic delays and their societal impact. While infrastructure limitations such as lane drops and sags are often blamed, driving behavior also plays a crucial role in shaping congestion dynamics. Differences in aggressiveness, following distance, and reaction time can destabilize traffic flow, leading to erratic trajectories, abrupt braking, and a higher risk of rear-end collisions (SWOV, 2022).

These behavioral factors directly affect the onset of traffic breakdown, which typically occurs when the incoming flow exceeds the effective capacity of the road. Instead of solely expanding physical infrastructure, improving behavioral coordination—such as maintaining shorter yet safe headways—offers a complementary path to mitigate congestion and enhance stability.

Given the prominent role of driving behavior in congestion, researchers have long turned to car-following (CF) models to understand the microscopic interactions between vehicles. CF models aim to capture the response of a following vehicle to the movements of its leader on a single lane. Traditional models (Brackstone and McDonald, 1999) primarily focus on quantifiable states such as speed, acceleration, and position. However, these models often struggle to account for the full spectrum of human driving behaviors, which are far more complex and varied than can be represented by a handful of parameters (Higgs and Abbas, 2015).

Recent advances in data collection and analysis have spurred the development of data-driven CF models. Unlike traditional approaches that rely on rigid mathematical formulations, data-driven models can learn directly from high-resolution car-following data. This enables them to capture nuanced driving behaviors—such as individual variations in reaction time, acceleration patterns, and headway preferences—that are pivotal in understanding and mitigating traffic breakdown (Papathanasopoulou and Antoniou, 2015). By better representing the complexity of real-world driving, these human-like models provide deeper insights into how different driving styles contribute to traffic oscillations and bottlenecks. This paper aims to propose a driving style-aware car-following model, classify driving styles, and investigate the impact of revealed driving styles on traffic flow.

This paper is structured as follows. Section II reviews prior work on driving style, conventional and data-driven car-following (CF) models, and breakdown mechanisms, identifying key research gaps. Sections III–V present the methodology, dataset processing, and the proposed style-aware CF model, followed by platoon-scale experiments (Section VI) and their results (Section VII), while Section VIII concludes with main findings and contributions.

## II. Literature Review

Congestion occurs when traffic demand exceeds road capacity, often triggered by incidents or bottlenecks. Traffic oscillations may arise from longitudinal instabilities (e.g., reaction or spacing errors) or lateral maneuvers, yet their growth and decay are largely independent of initial triggers (Zheng et al., 2011). Later studies attribute these dynamics to driving patterns rather than simple acceleration asymmetry, showing that certain behavioral tendencies—such as aggressiveness or timidity—amplify wave propagation (Laval and Leclercq, 2010). However, most prior research still classifies behavior in coarse terms (e.g. aggressive and timid) and relies primarily on analytical or model-based approaches (e.g., Newell's CF model), instead of experimental validation using model-driven simulations. Identifying distinct driving style groups and building up a model-driven experiment validation is therefore essential.

Driving style differs from driving behavior, which is context-dependent (Sagberg et al., 2015), and instead reflects a driver's longer-term preference—such as maintaining shorter

headways. Yet, there remains no consensus on how driving styles should be grouped. Clustering based on stated-preference surveys identifies multiple style categories (ranging from a few to several dozen) and shows that truck driving tends to be more homogeneous than passenger-car driving (Higgs and Abbas, 2015; Chen and Chen, 2019). However, trajectory-based analyses (Laval and Leclercq, 2010) often rely on coarse classifications (e.g., timid vs. aggressive), leaving a finer-grained characterization of driving styles unresolved.

Rather than extracting styles independently of the car-following (CF) context (Chen and Chen, 2019), our approach integrates style learning into the CF modeling process. The premise is that driving style can act as an external latent factor to enhance CF prediction, while improved CF accuracy, in turn, refines style discovery. Accordingly, the style extractor and CF model are trained jointly within a unified framework.

Data-driven CF modeling has been extensively explored over the past decades. Early models using MLPs (Simonelli et al., 2009; Colombaroni and Fusco, 2014) or fuzzy logic (Chong et al., 2013) demonstrated potential but lacked temporal awareness . Deep sequence models such as RNNs, LSTMs, and GRUs significantly improved predictive performance by capturing reaction delays, heterogeneity, and behavioral asymmetry (Zhou et al., 2017; Huang et al., 2018; Wang et al., 2019), while sequence-to-sequence designs extended prediction horizons (Ma and Qu, 2020). More recently, physics-informed hybrids have enhanced robustness and interpretability by constraining neural networks with CF principles (Mo et al., 2021; Geng et al., 2023). Nevertheless, two key challenges persist: driving-style clustering is typically performed outside model training, weakening the coupling between learned representations and behavioral grouping.

To address these gaps, we proposes a style-aware CF training framework. This jointly learns driving styles and car-following dynamics, enabling new style discovery. This will be used to link microscopic driving behavior to macroscopic traffic performance and energy efficiency.

## III. Methodology

The overall approach is that [first embedder, then .... etc, linking also the sections of this chapter]

### A. Style Embedder

Following Sagberg et al. (2015), the proposed approach defines driving style as a habitual and time-independent way of driving, in contrast to driving behavior, which reflects context-dependent reactions. We assume that style remains stable across time, whereas behavior may vary with traffic conditions.

In this study we propose a data-driven embedder model that directly learns latent driving style representations from trajectory data.

*1) Style-related Features*

Informed by human-factor studies, the following quantities are incorporated as style-related features: average time headway $\bar{h}$ (aggressiveness, proficiency), average reaction time $\bar{\tau}$ (aggressiveness, anticipation), their variabilities $\sigma_h$ and $\sigma_\tau$ (proficiency, distraction), and vehicle lengths $l^F$ and $l^L$ (truck–car differences). Among these, $h$ and $\tau$ are not directly observable in the CF dataset, therefore we need to extract them.

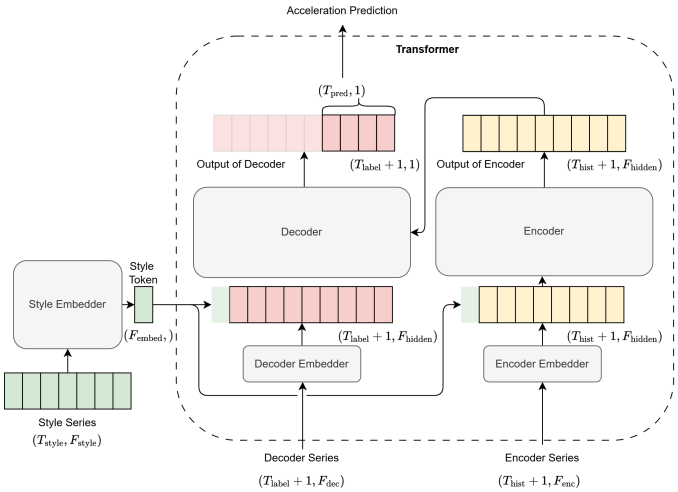


Fig. 1: Integrating the style embedding (token) into the Transformer Car-following model

Reaction time $\tau$ is estimated using a dynamic time warping (DTW) method, following the approach of Sharma et al. (2018). DTW is suitable for this task as it flexibly aligns leader and follower speed profiles, with the temporal offset reflecting the follower's response delay. However, when both vehicles are cruising at nearly the same speed, DTW may yield unrealistic results. To address this, an upper bound is imposed, $\tau \leq \Delta x/w$, where $w$ is fixed at -15 km/h.

Time headway $h$ is also included as a style-related feature, given its central role in classical car-following models.

Note that both $\tau$ and $h$ are extracted as time series and directly used as input features for the style embedder. For interpretability, however, their means and variances $(\bar{\tau}, \bar{h}, \sigma_\tau, \sigma_h)$ are employed to analyze the classification results and better characterize the identified styles.

*2) Architecture of Style Embedder*

The architecture builds on a Transformer encoder with an attention mechanism to preserve temporal dependencies such as reaction delays. As part of the style-aware CF framework (see Figure 1), the style embedder is jointly trained with the CF predictor, allowing it to benefit from the prediction loss and ultimately enhancing the overall accuracy of the car-following model.

The integration of the style embedding is shown in Figure 1, where the style token is prepended to both the encoder and decoder sequences and participates in training. This design allows the attention mechanism to dynamically decide whether and how to utilize the style information.

It outputs a fixed-length, time-independent style token that compactly represents driver traits while accommodating trajectories of arbitrary length. This embedding serves both as

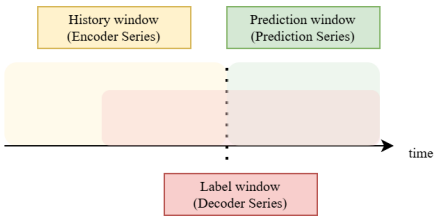


Fig. 2: Definitions of inputs and outputs of the Transformer model (also of the Transformer part in the Style-aware CF model)

input for downstream car-following models and as a basis for clustering analyses to examine whether consistent and interpretable driving style categories emerge.

### B. Style-aware CF Model

#### 1) Architecture

Based on the plain Transformer model, the Style-aware CF model introduces an additional style embedding. This token, generated by an independent Transformer encoder (the style embedder), encodes time-independent driving style information and is integrated with the encoder and decoder outputs to regulate the model's predictions.

As shown in Figure 1, the model takes three inputs: the encoder sequence, the decoder sequence, and a style series. The style series, originally temporal, is transformed by the style embedder into a time-independent token and prepended into the CF Transformer. This design enables the embedder to be trained jointly, producing two outputs: (i) the style embedding, applicable to clustering and behavioral analysis, and (ii) the predicted acceleration for car-following tasks. The attention mechanism adaptively determines how style information is utilized, which can later be inspected through attention matrices.

#### 2) Inputs and Outputs

Figure 2 depicts the segmentation of a single training sample. The historical window $T_{\text{hist}}$ and prediction window $T_{\text{pred}}$ are separated by a dashed line, while the label window $T_{\text{label}}$ spans both. The decoder can only access past information because the causal mask (Vaswani et al., 2017), hiding all future leader features during prediction. Conversely, follower features in $T_{\text{pred}}$ are replaced with historical averages within the $T_{\text{hist}}$ to prevent any leakage of future information.

The key difference from the plain CF Transformer is the additional style series input (Figure 1), constructed from features $(\tau_t, h_t, \Delta x_t, v_t^F)$ (Table I). While $\tau_t$ and $h_t$ represent driving style, $\Delta x_t$ and $v_t^F$ capture regime information. During training, $v_t^F$ in $T_{\text{pred}}$ is masked. A structured summary of inputs, outputs, and model configuration is provided in Table II.

### C. Model Training

#### 1) Details

The Style-aware CF model was trained on a workstation equipped with an Intel i7-12700K CPU, an NVIDIA RTX 3070Ti GPU (8 GB VRAM), and 32 GB RAM, using PyTorch 2.1 with CUDA 12.1. Training was conducted for 30 epochs with a batch size of 64. The Adam optimizer was employed with an initial learning rate of $10^{-4}$ and a weight decay of $10^{-5}$. Kinematic position loss (see Equation 1) was used as the training objective, and a step-based learning rate scheduler reduced the learning rate by a factor of 0.1 every 10 epochs. See Table II for hyper-parameter details.

On this setup, the full training required approximately 26 minutes. Model checkpoints were saved according to the lowest validation loss, and all experiments were run with a fixed random seed, 42 in this training, to ensure reproducibility. The training process demonstrated stable convergence, with validation loss plateauing after 25 epochs.

All code and scripts used for model training, evaluation, and simulation are publicly available at https://github.com/liheng423/style-cf.

#### 2) Loss Function

Following the kinematic model in Equation 1, position is obtained by double integration of acceleration, the model's output. The loss is defined on position accuracy, which implicitly requires consistent speed and acceleration predictions.

$$\begin{aligned} \hat{v}_t^{\mathrm{F}} &= v_0^F + \Delta t \sum_{i=1}^{T_{\text{pred}}} \hat{a}_i^F, \\ \hat{x}_t^{\mathrm{F}} &= x_0^F + \Delta t \sum_{i=1}^{T_{\text{pred}}} \hat{v}_i^F, \\ \text{Loss} &= \tfrac{1}{T_{\text{pred}}} \sum_{i=1}^{T_{\text{pred}}} \left(\hat{x}_i^F - x_i^F\right)^2, \end{aligned} \tag{1}$$

Here, $\hat{x}_t^F, \hat{v}_t^F, \hat{a}_t^F$ are predictions of followers, and $x_t^F, v_t^F, a_t^F$ are ground truth of followers, initialized from the state at the last step before prediction. This position-based MSE enforces physical consistency and prevents unrealistic accelerations, leading to stable and behaviorally plausible outputs.

### D. Model Testing Framework

#### 1) Extending to Long-duration Simulation

The Style Transformer and benchmark models are trained as short-duration predictors with horizons of only a few seconds (e.g., 4 s for the Style-aware CF). This design reduces computational cost, but strong short-duration performance may not generalize to longer horizons, where errors accumulate. Since we focuses on driving style and congestion formation, long-duration prediction offers a stringent test of model generalization and a robust basis for traffic breakdown replication.

To extend short-duration models to longer horizon, a recursive prediction framework is applied (Figure 3). The long-duration trajectory is generated by repeatedly shifting the prediction window forward by $T_{\text{pred}}$. The first history segment is initialized with ground truth, but subsequent windows increasingly rely on the model's own outputs. Given the 90-second sequence length, the effect of the initial ground-truth input is negligible in the long run.

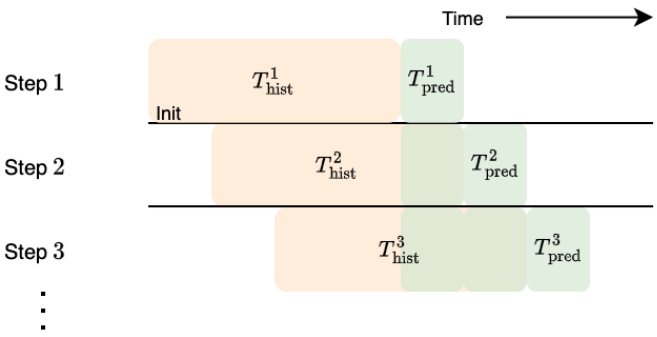


Fig. 3: Recursive prediction over time

TABLE I: Driving style-related car-following quantities

| Symbol | Quantity | Related human factors |
|---|---|---|
| $\bar{h}$ | Average time headway | Reflects aggressiveness and possibly driving proficiency |
| $\bar{\tau}$ | Average reaction time | Associated with aggressiveness, driving proficiency, and temporal anticipation |
| $\sigma_\tau$ | Standard deviation of reaction time | Related to driving proficiency and distraction, as variability may indicate inconsistent behavior over time |
| $\sigma_h$ | Standard deviation of time headway | Also associated with driving proficiency and distraction, similar to $\sigma_\tau$ |
| $l^F$ | Length of follower vehicle | Truck drivers exhibit different behavior compared to passenger car drivers (Ossen and Hoogendoorn, 2011) |
| $l^L$ | Length of leader vehicle | Driver behavior may differ when following a truck rather than a passenger car |

The car-following simulation algorithm employs recursive predictive modeling to derive the trajectory of a following vehicle, contingent upon the motion dynamics of the leading vehicle. This process incorporates a sliding history window alongside recursive input updates. The implementation of this simulation facilitates the assessment of performance over extended time periods.

After visualizing the clustering results, it is necessary to interpret the behavioral characteristics of each group. To this end, Table I summarizes the style-related CF indicators, which are constructed based on human factors in car-following (Sagberg et al., 2015). These quantities capture aggressiveness, reaction consistency, distraction, and the influence of vehicle types, providing a meaningful basis for distinguishing driving styles across clusters. To verify that these differences reflect genuine behavioral variations rather than artifacts of the learned embedding space, token-replacement experiments are performed: the style and vehicle-type tokens of each sample are swapped with those from other groups, and the resulting behavioral changes are compared.

*2) Incorporating Style in Testing*

Different models require different input structures, so a consistent evaluation setup is needed to ensure fair comparison. For non–style-aware models (LSTM, plain Transformer, and IDM), explicit style information is not used; instead, average-style representations are adopted. In particular, the IDM is calibrated with the mean parameter values from all training trajectories (Table III). For the Style Transformer, style information is incorporated through group-specific average embeddings. To obtain these, 30% of the long test trajectories are used for style learning, as their 90-second duration yields more stable representations. The data are evenly sampled across groups to avoid bias. Meanwhile, clustering is performed using the learned embedder and k-means on the training set, which contains non-overlapping trajectories from the same set of vehicles. Since the training dataset covers all vehicle IDs in the test set, no trajectory lacks a style label. The final group-average embeddings are then directly injected into the model, bypassing the embedder during inference. The entire process is summarized in Figure 4.

### *E. Benchmark Models*

Three models are considered in this study: the rule-based Intelligent Driver Model (IDM), a Long Short-duration Memory (LSTM) network, and a Transformer-based model. An overview of the hyperparameters, along with the model inputs and outputs, is provided in Table II.

The IDM serves (Treiber et al., 2000) as a baseline, simulating acceleration based on current velocity, spacing, and leader speed. To ensure fair comparison, the IDM is calibrated on the same training dataset using a genetic algorithm (Mitchell, 1996), which estimates parameters per car-following pair. For benchmarking, the IDM is constrained to uniform parameters across all pairs. Therefore, average values from the calibrated distribution are adopted as the final setting (Table III). Calibration is performed on 1000 randomly sampled trajectories using the Genetic Algorithms.

The LSTM model (Hochreiter and Schmidhuber, 1997) uses a 6-second historical window with features including vehicle velocities, spacing, and lengths to account for heterogeneous vehicle types. It predicts a sequence of accelerations but cannot access future leader behavior due to the sequential structure, limiting its long-duration prediction ability.

The Transformer model (Vaswani et al., 2017) extends this by incorporating future leader behavior through a label window while masking follower movements to prevent data leakage. It enables longer prediction horizons and better handles long-range dependencies, though it requires careful input segmentation and masking design.

## IV. Data Used

### *A. Data Source and Requirements*

For the study we use the Zen traffic dataset, providing vehicle trajectories from freeways. It includes free-flow and congestion states. We remove trajectories passing through

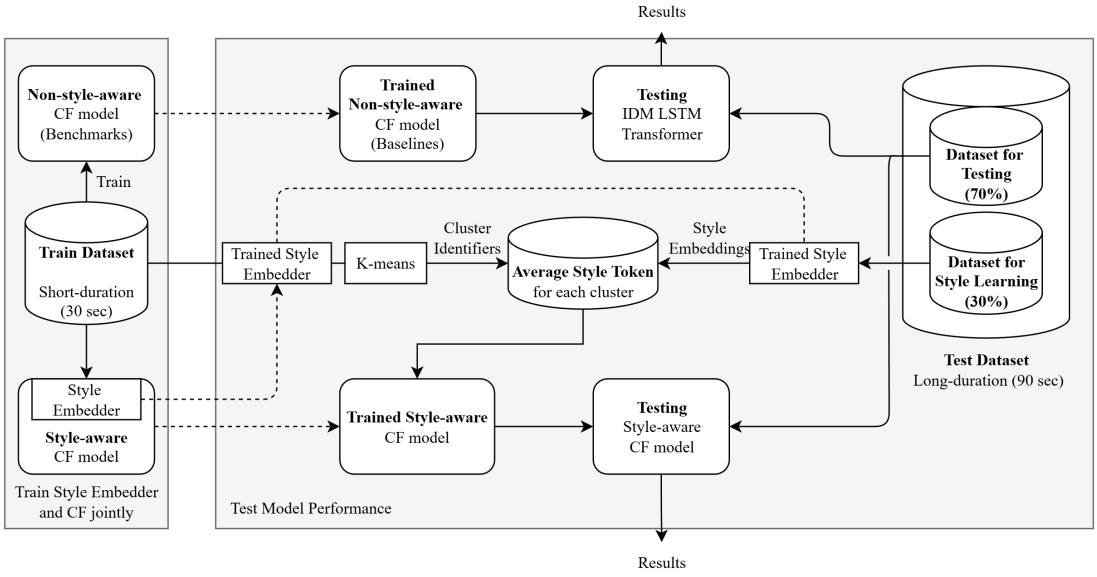


Fig. 4: Training and testing data handling involving style information

TABLE II: Detailed parameters in training process ($\langle v^F \rangle_{T_{\text{label}}}$ indicates the average value of $v^F$ in the label window)

| Param. | LSTM | Transformer | Style-aware CF | IDM |
|---|---|---|---|---|
| **Hyper-parameter Setting** | | | | |
| LR | 0.001 | 0.001 | 0.001 | - |
| Batch | 64 | 64 | 64 | - |
| Epochs | 20 | 30 | 30 | - |
| Dropout | 0.0 | 0.1 | 0.1 | - |
| Activation | Sigmoid | ReLU | ReLU | - |
| MaxNorm | 10.0 | 10.0 | 10.0 | - |
| Optimizer | Adam | Adam | Adam | - |
| **Data Segmentation Setting** (sec) | | | | |
| $T_{\text{pred}}$ | 2 | 4 | 4 | 0.1 |
| $T_{\text{hist}}$ | 6 | 6 | 6 | 0.1 |
| $T_{\text{label}}$ | - | 4 + 4 | 4 + 4 | - |
| $T_{\text{hist}}$ | - | - | 6 | - |
| **Features** | | | | |
| History | $v^F, \Delta v, \Delta x, l^F, l^L$ | | | $v^F, v^L, \Delta x$ |
| $\text{Label}_{\text{hist}}$ | - | $v^L, v^F$ | | - |
| $\text{Label}_{\text{pred}}$ | - | $v^L, \langle v^F \rangle_{T_{\text{label}}}$ | | - |
| Style | - | - | $(\tau, h, \Delta x, v^F)$ | - |
| Output | $a^F$ | $a^F$ | $a^F$ | $a^F$ |

TABLE III: Average calibrated IDM parameters (Zen dataset)

| $v_0^F$ | $\Delta x_0$ | $h_0$ | $a_{\max}$ | $b$ |
|---|---|---|---|---|
| 24.70 m/s | 1.70 m | 1.19 s | 1.70 m/s$^2$ | 2.53 m/s$^2$ |

weaving areas or changing lane and select only pairs that have a constant leader throughout the analysis time.

Before usage, we enforce enforce kinematic consistency of the data. using Kalman filter.

Subsequently, we perform validity checks on headway, terminal position error, and acceleration to remove physically implausible samples. These procedures yield a clean and dynamically consistent set of car-following segments that serve as the foundation for further analysis.

In this section, we instead examine the distribution of vehicle types in the dataset to demonstrate that the model can generalize its learned behavior to truck-following cases during testing (see Table IV)

*1) Car-following Model Dataset*

**Train & Validation Dataset (Short-duration CF segments)**: The processed short-duration dataset contains 44,207 car-following pairs, with 38,176 passenger-car ($l < 7.5$m) followers and 6,031 truck followers ($l > 7.5$m), thus preserving vehicle heterogeneity. Both velocity and acceleration distributions fall within realistic ranges, consistent with Montanino and Punzo (2015).

TABLE IV: Number of vehicle types in the processed datasets (short- and long-duration CF segments, unit [pair])

| **Category** | **Train** | **Test** |
|---|---|---|
| Truck–Truck | 1,786 | 84 |
| Truck–Passenger Car | 7,944 | 686 |
| Passenger Car–Truck | 4,245 | 194 |
| Passenger Car–Passenger Car | 30,232 | 3,409 |
| **Total (Truck Followers)** | 6,031 | 278 |
| **Total (Passenger-car Followers)** | 38,176 | 4,095 |
| **Grand Total** | 44,207 | 4,373 |

**Test Dataset (Long-duration CF segments)** : The test dataset is constructed with a 90-second time window, yielding 2,640 samples after filtering. In terms of its speed distribution, it covers traffic states from congestion to free flow, ensuring consistency with the training set. The vehicle composition (2,519 passenger-car followers and 121 truck followers)

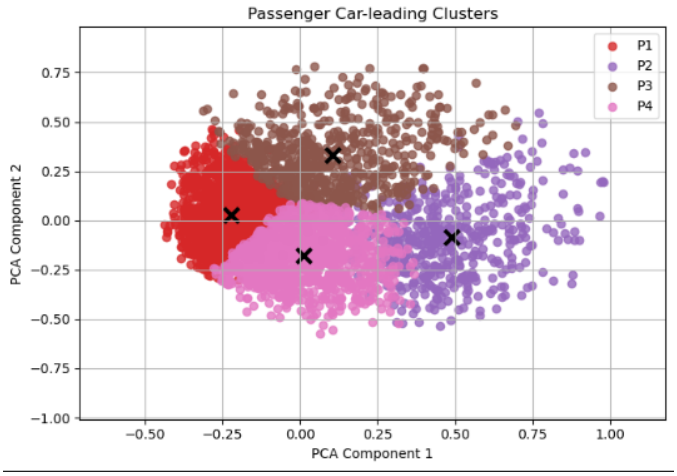


Fig. 5: Clustering of driving style embeddings

TABLE V: Cluster-wise summary statistics

| | P1 | P2 | P3 | P4 |
|---|---|---|---|---|
| $\bar{\tau}$ | 1.77 ± 0.39 | 2.00 ± 0.42 | 2.18 ± 0.40 | 1.71 ± 0.40 |
| $\bar{h}$ | 1.31 ± 0.33 | 1.43 ± 0.28 | 1.65 ± 0.22 | 1.12 ± 0.30 |
| $\Delta\bar{x}$ | 15.11 ± 4.87 | 19.73 ± 5.30 | 21.14 ± 5.06 | 14.63 ± 4.66 |
| $\bar{v}^F$ | 11.68 ± 2.69 | 13.68 ± 2.75 | 12.76 ± 2.68 | 13.14 ± 2.84 |
| $\sigma_\tau$ | 1.01 ± 0.19 | 1.14 ± 0.17 | 1.14 ± 0.15 | 1.04 ± 0.20 |
| $\sigma_h$ | 0.28 ± 0.12 | 0.40 ± 0.44 | 0.40 ± 0.17 | 0.30 ± 0.09 |
| Count | 1187 | 383 | 508 | 915 |

largely matches the training data, though truck-follower samples remain limited due to the short trajectory duration.

## V. Embedding and CF Model Evaluation

Based on the trained CF model and the style-embedder, this section evaluates two aspects: (i) the quality of the clustering and the influence of the extracted style embeddings, and (ii) the joint performance of the CF model with the style embedder to assess how these embeddings enhance prediction accuracy.

### A. Driving Style Clustering

The style embedding (token) is time-independent, allowing direct clustering in the embedding space. To avoid data leakage, clustering is conducted on the long-duration dataset, which is independent of training/validation data and contains 2519 passenger-car follower samples.

K-means is used to partition embedding vectors into $K$ clusters by minimizing within-cluster variance. In this study, it groups segments with similar style representations, thereby identifying distinct driving patterns.

Table V shows the properties of the 4 found clusters. They are visualized in Figure 5 in two dimensions. P1 and P4 appear compact and dominant, suggesting stable styles, whereas P2 and P3 show greater variability. Overall, results indicate that driving style lies on a behavioral spectrum rather than discrete categories, consistent with Sagberg et al. (2015).

The clusters reveal distinct behavioral patterns (Table VI). P1 shows short headway and small spacing, suggesting steady yet impatient driving. P2 has longer reaction and high variability, reflecting occasional hesitation. P3 features the longest headway and spacing, consistent with defensive and conservative driving. P4 reacts fastest but with instability, aligning with aggressive and impulsive behavior.

TABLE VI: Interpretation of clusters

| Cluster | Style | Key features | Human factors |
|---|---|---|---|
| P1 | Stable-aggressive | Short $\bar{h}$, small spacing, low $\sigma_h$ | Mild aggressiveness, speed focus |
| P2 | Unstable-timid | Longer $\bar{\tau}$, higher $\sigma_h$ | Occasional hesitation, cautious balance |
| P3 | Stable-timid | Longest $\bar{\tau}$, largest spacing | Most cautious, safety-oriented |
| P4 | Unstable-aggressive | Lowest $\bar{h}$, lowest $\bar{\tau}$, higher $\sigma_h$ and $\sigma_\tau$ compared to P1 | Aggressive, unstable control |

### B. Driving Style Embedding Validation

The style embedder is trained jointly with the car-following model, enabling the extraction of representations that improve prediction accuracy. We now validate whether the learned embeddings are functionally meaningful and behaviorally interpretable. Three complementary tests are conducted.

*1) Truck as Follower*

Since truck samples are scarce in the dataset, we emulate them by setting follower length $l^F = 10$ m. As shown in Figure 6, this modification consistently yields shorter reaction times while time headway remains largely unchanged. This indicates that trucks at low speeds do not follow differently in terms of headway but react faster, consistent with Ossen and Hoogendoorn (2011), which links such behavior to professional training.

*2) Cross-over Test: Effectiveness of the Embedding*

To separate driving style from scene effects, we replace embeddings with cluster centroids (Figure 7). Each row corresponds to the original cluster, and each column to the substituted embedding. Shifts in time headway $\bar{h}$ across substitutions confirm that embeddings encode behavioral traits independent of context. In particular, P2 and P3 induce the strongest changes, consistent with their interpretation as conservative and aggressive styles (Table VI). These results demonstrate that embeddings systematically modulate car-following behavior.

### C. Model Testing Results

The long-duration prediction task provides a common evaluation framework across all models. To assess performance, we report multiple error metrics, their temporal evolution, example trajectories, and attention weights.

*1) Error Statistics*

Prediction performance is evaluated using MSE, MAE, and RMSE of follower position $x_t^F$, computed per trajectory. This choice reflects the loss design (Equation 1) and is widely used in CF studies (Huang et al., 2018). In addition to mean and standard deviation, quartiles (Q1–Q3) are reported to capture

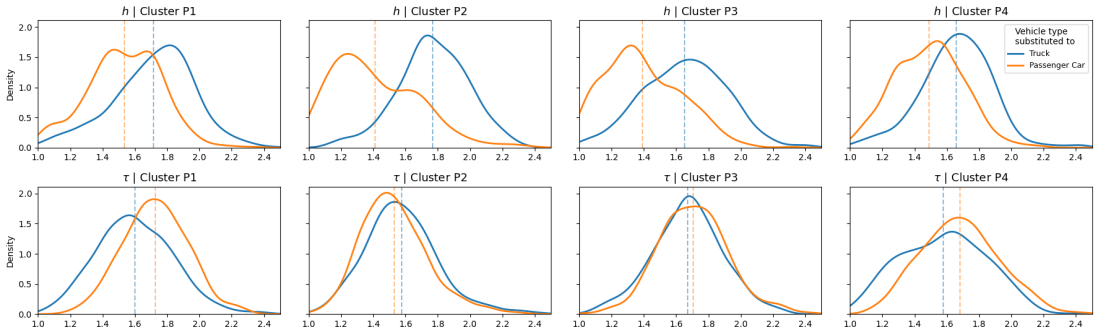


Fig. 6: Truck emulation ($l^F = 10$ m) across clusters

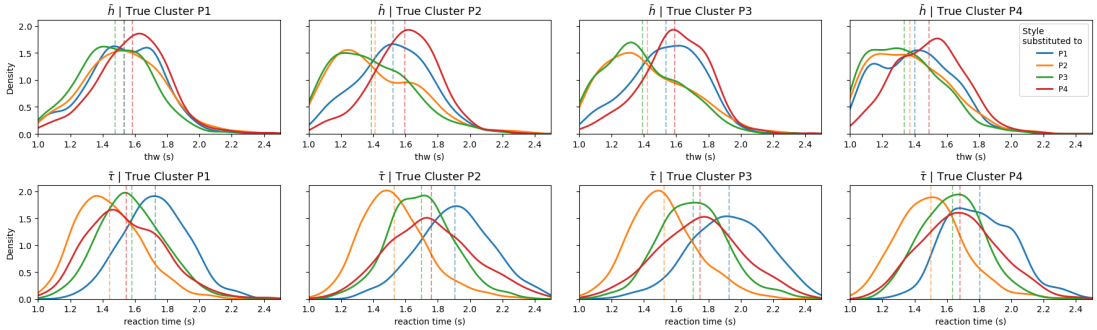


Fig. 7: Embedding substitution validation (cluster centroids)

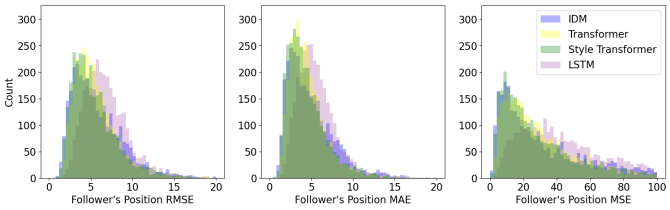


Fig. 8: Distribution of position error metrics

distributional properties under different traffic regimes. Results are summarized in Table VII, and distributions are shown in Figure 8. Overall, the Style-aware CF achieves the lowest average errors, while LSTM performs worst.

TABLE VII: Performance statistics of different models on prediction metrics (rounded to 2 decimals).

| Metric | Statistic | Style-aware | Trans. | LSTM | IDM |
|---|---|---|---|---|---|
| MSE | Mean | 33.71 | 32.71 | 53.33 | 41.16 |
| | Std Dev | 41.18 | 33.89 | 47.65 | 48.79 |
| RMSE | Mean | **5.17** | 5.21 | 6.79 | 5.64 |
| | Std Dev | 2.68 | 2.35 | 2.69 | 3.06 |
| MAE | Mean | **4.24** | 4.27 | 5.58 | 4.83 |
| | Std Dev | 2.35 | 2.03 | 2.33 | 2.84 |

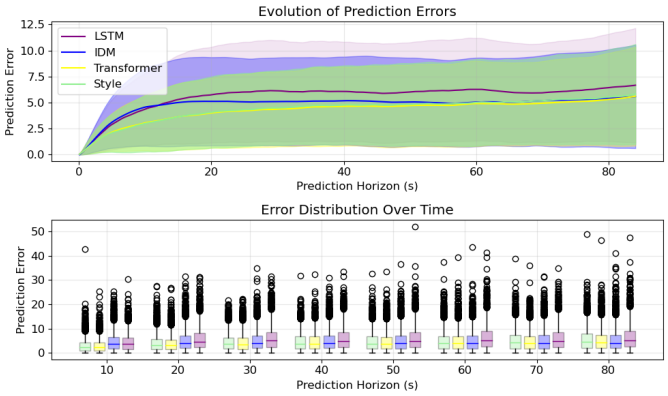


Fig. 9: Evolution of position MAE over time

*2) Error Evolution*

Recursive prediction causes error accumulation. Figure 9 shows that all models exhibit increasing MAE with time, but LSTM diverges fastest and has the largest variance. In contrast, the Style Transformer and IDM remain more stable, reflecting the benefit of incorporating style information—either through tokens or calibrated parameters.

*3) Example Trajectories*

Two 90-second test cases illustrate model behavior (Figure 10). IDM responds without delay due to its memoryless design, while the Transformer-based models capture realistic reaction delays. Style Transformer and IDM, both guided by style information, show more consistent spacing and smoother

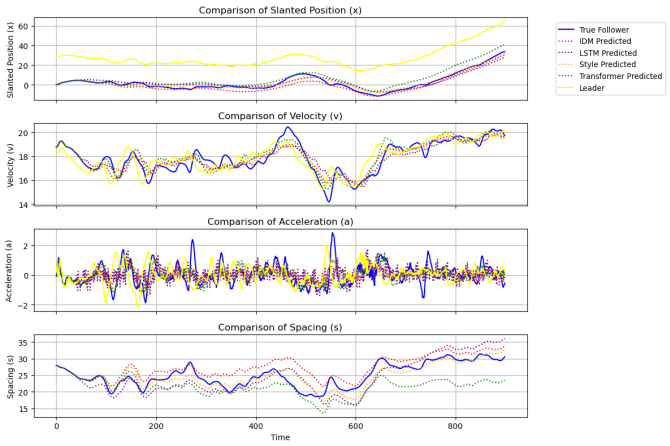


Fig. 10: Trajectory comparison (Sample 200, slanted term: $-17.5t$)

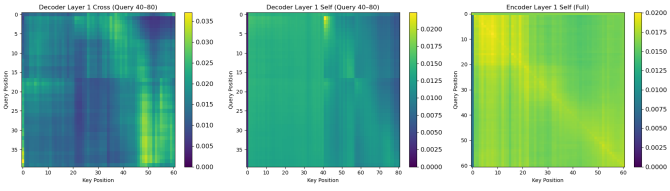


Fig. 11: Attention matrix example (Sample 233)

accelerations, whereas the plain Transformer and LSTM are more volatile and sensitive to short-term fluctuations.

*4) Attention Matrix*

Figure 11 shows attention maps, where the y-axis denotes query positions (output) and the x-axis lists key positions (input) starting with the prepended style token followed by input steps (1–6 s). In the decoder cross-attention, the style token is almost ignored for near-term queries ($<$ 2 s) but becomes strongly attended after roughly 2–2.5 s, providing global guidance for longer-horizon prediction, while decoder self-attention remains diffuse yet strictly causal, using only past and already generated steps within the prediction window. By contrast, encoder self-attention exhibits a sharp autoregressive diagonal over the entire history, emphasizing its role in encoding strong temporal dependencies and preserving sequential structure.

## VI. Platoon Experiment Setups

The following sections investigate how different driving styles, alone and in combination, influence platoon dynamics and breakdown, based on the unified car-following model and four style categories introduced earlier.

### A. Experiment Framework

*1) Overview*

We extend the simulator to the platoon scale, where outcomes depend on both (i) the leader trajectory and (ii) platoon composition. Leader trajectories serve as control variables, while platoon composition—including style mix, vehicle type, and ordering—determines collective response.

a) Setup: Define platoon composition and select a leader trajectory covering diverse regimes.
b) Assemble: Sample vehicles by style class and assign the chosen leader.
c) Simulate: Run for 240 s and record outputs.
d) Analyze: Compute indicators and assess results.

### B. Metrics for Traffic Breakdown

In this study, multiple indicators are adopted to describe the magnitude of the traffic breakdown and the effect of driving style on traffic flow (Table VIII). The following sections introduce their calculation and characteristics.

| Indicator | Unit | Sig. Test | Description |
|---|---|---|---|
| Delay | [s] | two-way ANOVA | Time lost relative to Newell baseline |
| ST | [s] | two-way ANOVA | Total time with $v < 1$ m/s, centered per leader |
| FC | [L/km] | two-way ANOVA | Fuel used per kilometer |
| $w_{\text{acc}}$ | [km/h] | t-test | Speed of accelerating waves |
| $w_{\text{dec}}$ | [km/h] | t-test | Speed of decelerating waves |
| $q_{\max}$ | [veh/h] | WLS t-test | Capacity from fitted FD |
| $k_{\text{jam}}$ | [veh/km] | WLS t-test | Jam density from congested intercept |
| $k_{\text{crit}}$ | [veh/km] | WLS t-test | Critical density (triangular FD) |
| $b$ | [km/h] | WLS t-test | Slope of congested branch (backward wave speed) |

TABLE VIII: Indicators employed in the platoon experiments (ANOVA: analysis of variance, WLS: weighted least squares)

*1) Trajectory-level Measures*

To evaluate the overall performance of a platoon, the following indicators are designed. Note, for each indicator, two-way ANOVA tests overall differences across style groups.

**Delay** is defined as the excess travel time relative to a disturbance-free benchmark (Newell's model):

$$\text{Delay} = TTS_{\text{actual}} - TTS_{\text{benchmark}}, \tag{2}$$

where $TTS$ is computed as the cumulative time vehicles remain upstream of the boundary $L$.

**Standstill time (ST)** measures immobility as the cumulative time vehicles operate in the standstill regime:

$$\text{ST} = \Delta t \cdot \sum_{n=1}^{N} \sum_{k=0}^{K-1} \mathbf{1}_{\{v_n^k < 1 \text{ m/s},\, -0.1 < a_n^k < 0.1\}}, \tag{3}$$

with $v_n^k$ and $a_n^k$ denoting the speed and acceleration of vehicle $n$ at time step $k$.

**Fuel Consumption (FC)**: Traffic breakdowns increase energy costs due to stop-and-go patterns, which cause frequent acceleration and deceleration, raising fuel consumption and environmental impact. Evaluating energy efficiency with breakdown dynamics connects operational stability to sustainability. Fuel consumption is estimated using the VT-Micro model, which predicts rates based on speed $v(t)$ and acceleration $a(t)$ with the equation:

$$J_{\text{FC}}(t) = \exp\left(\tilde{v}(t)^{\top} P_{\text{FC}} \tilde{a}(t)\right),$$

where $\tilde{v}(t) = [1, v, v^2, v^3]^\top$ and $\tilde{a}(t) = [1, a, a^2, a^3]^\top$. This model, using calibrated coefficients from Zegeye et al. (2013), effectively captures stop-and-go dynamics by distinguishing acceleration regimes.

**Statistical Testing** is usde to compare indicator values fairly across groups. The analysis must control for differences in leader movements, since the leader's trajectory affects follower behavior. We therefore apply a two-way ANOVA without interaction terms, written as a linear model with treatment (reference) coding:

$$y_{p\ell i} = \mu + \sum_{p \neq p_0} \alpha_p D_{ip} + \sum_{\ell \neq \ell_0} \gamma_\ell D_{i\ell} + \varepsilon_{p\ell i}, \tag{4}$$

where $y_{p\ell i}$ is the response in platoon $p$ under leader environment $\ell$. The dummy variables $D_{ip}$ and $D_{i\ell}$ capture mean differences relative to the reference platoon $p_0$ and reference leader $\ell_0$. Thus, $\alpha_p$ measures how the indicator in platoon $p$ differs from the reference after accounting for leader effects, and $\gamma_\ell$ reflects the shift of indicator introduced by leader $\ell$. Significance is evaluated using the associated $t$-tests, consistent with the standard two-way ANOVA framework.

*2) Traffic Waves*

Traffic waves, propagating from downstream to upstream, mark transitions between free-flow and congestion. A decelerating wave signals breakdown onset, while an accelerating wave reflects recovery. This study focuses on wave speed—the upstream propagation velocity—since it is closely tied to driver behavior (Chen et al., 2014). Other characteristics (e.g., duration) are excluded, as trained agents ensure waves consistently propagate upstream without premature dissipation.

Wave detection proceeds in three steps: (i) **Event detection** — each velocity series $\tilde{v}_i(t)$ is smoothed over a window $W_s$ to suppress noise, and regime-change events are extracted as local extrema within a look-around window $W_\ell$; (ii) **Pairwise matching** — leader–follower events within $\tau_{\max}$ are linked, yielding temporal lag $\Delta t$, spatial gap $\Delta x$, and wave speed $c = \Delta x / \Delta t$; (iii) **Chaining** — valid links are combined into sequences spanning at least $L_{\min}$ vehicles, with wave speed summarized by the median $c$.

TABLE IX: Parameter settings for wave detection

| Parameter | $W_\ell$ | $L_{\min}$ | $\tau_{\max}$ | $W_s$ |
|---|---|---|---|---|
| Value | 1.0 s | 5 | 1.0 s | 2.0 s |

The algorithm outputs two categories of speeds: decelerating waves $w_{\text{dec}}$ (breakdown) and accelerating waves $w_{\text{acc}}$ (recovery). These indicators will later serve to distinguish the traffic wave speeds generated by different driving styles.

To assess whether wave speeds differ across groups, pairwise comparisons are performed using two-sample t-tests, where each group is tested against a chosen reference baseline. Normality of the distributions is checked beforehand, since the data come from different leaders and may not perfectly follow a normal distribution.

*3) Macroscopic Variables*

Macroscopic traffic variables—flow, density, and speed—characterize platoon conditions in a given time–space region. Following Edie's definition (Edie, 1965), they can be consistently retrieved from any arbitrary region.

*a) Fundamental Diagram (FD) Construction*

To capture platoon dynamics, parallelograms are used as measurement units (Figure 12a), covering full trajectories while minimizing void areas.

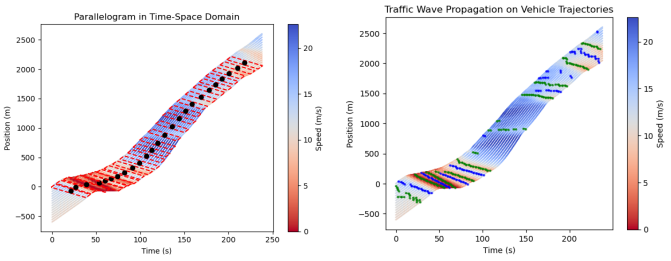


(a) Parallelogram generation (b) blue: dec., green: acc.

Fig. 12: Parallelogram construction (Left) and wave detection (Right) (Sample 16).

The congested branch of the fundamental diagram is approximated by a linear fit

$$q(k) = a + bk, \qquad w = -b, \tag{5}$$

where $w$ denotes the backward wave speed. Under a triangular FD assumption, jam density, critical density, and capacity follow as

$$k_j = \frac{a}{w}, \qquad k_{\text{crit}} = \frac{a}{v_f + w}, \qquad q_{\max} = v_f \, k_{\text{crit}}, \tag{6}$$

with $v_f$ given by the 90th percentile of the leading-vehicle speed (fixed to 50.4 km/h).

*b) Statistical Inference of FD Parameters*

Weighted Least Square (WLS) estimates all four style groups in one model, using a clear reference category. The congested branch of the fundamental diagram is approximated by a linear function, making least squares estimation suitable. Since density–flow data is concentrated near the critical density and sparse at higher densities, a weighting scheme emphasizes high-density observations to prevent mid-density bias.

To assess differences across platoons, a WLS model with density–platoon interactions is estimated:

$$q_i = \beta_0 + \beta_1 (k_i - k_0) + \sum_{s=2}^{S} \gamma_s D_{is} + \sum_{s=2}^{S} \delta_s (k_i - k_0) D_{is} + \varepsilon_i, \tag{7}$$

where $\beta_1$ is the baseline slope and $\delta_s$ represents deviations for platoon $s$. And $D_{is}$ is the dummy variable. T-tests are performed for pairwise comparisons to the given baseline group.

For each platoon $s$, parameters are obtained as

$$a_s = (\beta_0 - \beta_1 k_0) + \gamma_s - \delta_s k_0, \qquad b_s = \beta_1 + \delta_s, \qquad w_s = -b_s, \tag{8}$$

and the FD quantities $(k_{j,s}, k_{\mathrm{crit},s}, q_{\max,s})$ are derived by substituting $(a_s, w_s, v_{f,s})$ into the general formulas above. Here $v_{f,s}$ is taken as the 90th percentile of the leading-vehicle speed (50.4 km/h), and bootstrap resampling provides confidence intervals for style-wise differences.

### C. Experiment Setup

To investigate the influence of driving style on traffic breakdown, a series of platoon simulations were designed. Building on the classification results (see Table VI), which identified groups with aggressive, timid, stable, and unstable characteristics, we directly incorporate these styles into the experimental setup. Given the large number of possible style permutations, the study focuses on representative compositions summarized in Table X.

Each simulation is independent, with the environment reset after completion.

Notations in Table X are clarified as follows:

- $T$ and $P$ denote trucks (9 m) and passenger cars (5 m) with the same driving style, randomly assigned in each run.
- Square brackets $[\cdot]$ indicate ordered lists.
- Curly braces $\{\cdot\}$ denote unordered sets specifying only style proportions.
- A subscript gives the platoon size.
- Percentages indicate composition shares.

TABLE X: Experiment set-up summary (each experiment is carried out with 15 replications and 10 different leaders)

| Experiment | Composition | Category |
|---|---|---|
| 1-a | $[P1, P1\cdots]^{16}$ | aggressive and stable |
| 1-b | $[P2, P2\cdots]^{16}$ | timid but unstable |
| 1-c | $[P3, P3\cdots]^{16}$ | timid and stable |
| 1-d | $[P4, P4\cdots]^{16}$ | aggressive but unstable |
| 2-a | $[P1, P2, P1, P2, \ldots]^{16}$ | mixed stability |
| 2-b | $[P1, P3, P1, P3\cdots]^{16}$ | stable |
| 2-c | $[P1, P4, P1, P4\cdots]^{16}$ | mixed stability |
| 2-d | $[P2, P3, P2, P3\cdots]^{16}$ | mixed stability |
| 2-e | $[P2, P4, P2, P4\cdots]^{16}$ | unstable |
| 2-f | $[P3, P4, P3, P4\cdots]^{16}$ | mixed stability |
| 2-b-1 | $\{P1(25\%), P3(75\%)\}^{16}$ | stable (P3 biased) |
| 2-b-2 | $\{P1(50\%), P3(50\%)\}^{16}$ | stable (balanced) |
| 2-b-3 | $\{P1(75\%), P3(25\%)\}^{16}$ | stable (P1 biased) |
| 3-a | $[T1, T1\cdots]^{16}$ | aggressive and stable |
| 3-b | $[T2, T2\cdots]^{16}$ | timid but unstable |
| 3-c | $[T3, T3\cdots]^{16}$ | timid and stable |
| 3-d | $[T4, T4\cdots]^{16}$ | aggressive but unstable |
| 3-e | $[P,\cdots]^4[T,\cdots]^4[P,\cdots]^4$ | clustered trucks at center |
| 3-f | $[T, T\cdots]^8[P, P\cdots]^8$ | clustered trucks as leader |
| 3-g | $[T, P, T, P\cdots]^{16}$ | alternating |
| 3-h | $\{T(50\%), P(50\%)\}^{16}$ | random |

Accordingly, three experiments are summarized as follows:

- **Experiment 1: Style Homogeneous Platoon** *Do timid drivers destabilize traffic more than aggressive ones?* Homogeneous platoons (P1–P4) are tested to isolate style-specific impacts, allowing systematic comparison of aggressive vs. timid and stable vs. unstable groups.
- **Experiment 2: Alternating vs. Homogeneous** *Is mixing styles more efficient?* Alternating platoons (e.g., P1–P3) are compared with homogeneous baselines. The most promising pair (P1–P3) is further tested under different proportions (25%, 50%, 75%).
- **Experiment 3: Truck Placement** *How does truck distribution affect efficiency?* Trucks are placed in clustered, alternating, or random configurations within passenger-car platoons to evaluate their impact on overall delay.

## VII. Platoon Experiment Results

### A. Experiment 1 - Style Homogeneous Platoon

#### 1) Trajectory level Indicators

TABLE XI: Difference of indicators compared to the reference level (1-c), * indicates the difference compared to the reference level (in bold text) significant ($p < 0.001$)

| Experiment | $\lvert w_{\mathrm{dec}}\rvert$ (m/s) | $\lvert w_{\mathrm{acc}}\rvert$ (m/s) | ST (S) | Delay (s) | FC (L/km) |
|---|---|---|---|---|---|
| | t-test | | WLS + t-test | | |
| 1-a | +0.55* | +0.56* | -26* | +11* | -0.48* |
| 1-b | +0.51* | +0.22* | +77* | +45* | +0.43* |
| **1-c** | **5.68** | **5.29** | **228** | **835** | **0.69** |
| 1-d | +1.01* | +1.05* | +38* | +41* | -0.45* |

Delay and standstill time provide complementary views of breakdown severity. Table XI summarizes the results for homogeneous platoons.

Aggressive groups (1-a, 1-d) show faster wave propagation and shorter standstills, indicating that although breakdowns emerge quickly, they also dissipate rapidly. By contrast, the conservative–unstable group (1-b) exhibits slower waves but much longer standstills, leading to heavier congestion.

From centered delay, the conservative–stable platoon (1-c) performs best, achieving significantly lower delay than all other groups. This highlights that long standstills do not necessarily worsen overall performance when stability mitigates breakdown propagation.

Fuel consumption reveals a different picture: conservative groups (1-b, 1-c) generally consume more fuel due to extended low-speed gliding. Notably, 1-c has the lowest delay but the highest fuel use, while 1-d demonstrates that high standstill time does not necessarily imply high fuel consumption, as rapid recovery to efficient cruising offsets prolonged stops.

#### 2) Detected Traffic Wave

Traffic waves differ substantially across driving styles, offering key insights into how breakdowns form and dissipate. Unlike wave speeds inferred from fundamental diagrams, which rely on macroscopic calibration, here the waves are directly detected from trajectories and classified into accelerating and decelerating types—an essential distinction for breakdown dynamics.

Figure 13 shows mean accelerating and decelerating wave speeds with 95% confidence intervals (bootstrapped). Three findings emerge:

a) **Near symmetry:** points cluster around the line $y = x$, with accelerating waves slightly faster in magnitude.

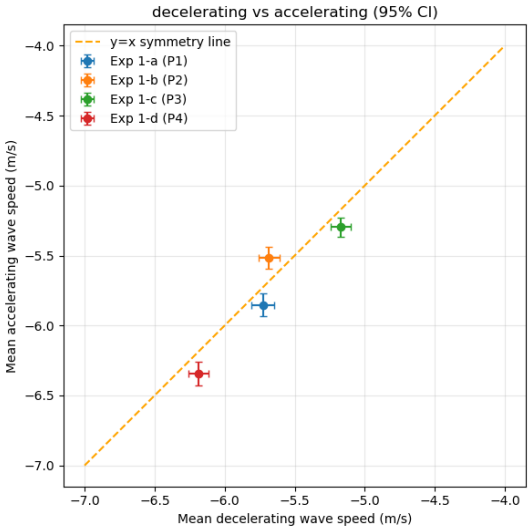


Fig. 13: Accelerating vs. decelerating wave speeds (95% CI).

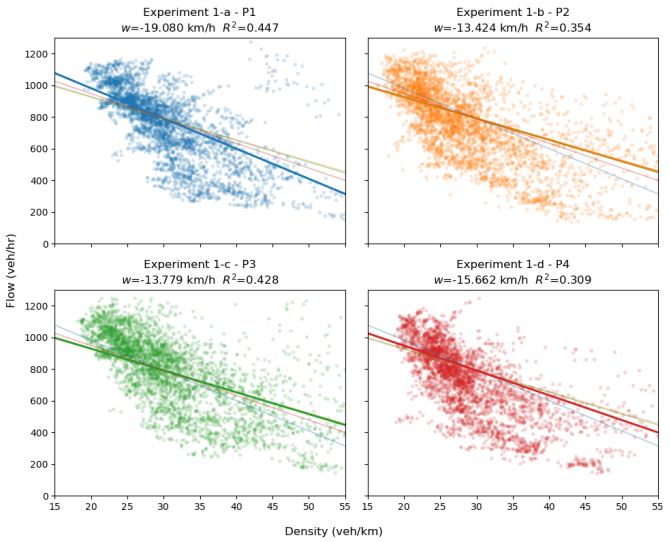


Fig. 14: Calibrated Fundamental Diagrams (Congested Branch)

b) **Style effect:** Group 1-d (*aggressive–unstable*) has the largest wave speeds, while Group 1-c (*conservative–stable*) shows the smallest.

c) **Fragility:** Group 1-b (*conservative-unstable*)'s confidence interval approaches or crosses $y = x$, implying possible inversion (accelerating slower than decelerating), signaling greater risk of persistent breakdown.

*3) Macroscopic Traffic Variables*

Macroscopic traffic variables, obtained via fundamental-diagram calibration, provide an aggregate view of style-dependent differences (Table XII, Figure 14). Only statistically significant contrasts are highlighted below.

Four main observations emerge:

a) **Wave speed:** Steeper slopes in 1-a (*aggressive–stable*) and 1-d (*aggressive–unstable*) indicate faster reactions, while gentler slopes in 1-b (*timid–unstable*) and 1-c (*timid–stable*) reflect slower responses, consistent with Table XI.

b) **Jam density:** Defensive 1-c tolerates the highest densities, whereas unstable 1-d sustains the lowest.

c) **Maximum flow:** Efficiency–stability trade-offs arise: 1-a yields the highest flow, 1-b the lowest, 1-c a moderate but reliable flow, and 1-d's instability reduces its advantage.

d) **Critical density:** Stable groups—1-a and 1-c—maintain significantly higher $k_{\text{crit}}$, making them less prone to breakdown near capacity.

TABLE XII: Fundamental-diagram parameters estimated from WLS (reference = 1-c). Significance indicates difference from the reference group( * denotes $p < 0.001$).

| Exp. | $a$ (veh/hr) | $b$ (km/hr) | $k_{\text{jam}}$ (veh/km) | $k_{\text{crit}}$ (veh/km) | $q_{\max}$ (veh/hr) |
|---|---|---|---|---|---|
| | WLS + t-test | | Bootstrap difference testing | | |
| 1-a | 1493.38* | -24.34* | 61.36* | 19.06* | 1029.40* |
| 1-b | 1294.70 | -18.73 | 69.11 | 17.80 | 961.25 |
| **1-c** | **1393.23** | **-19.52** | **71.37** | **18.95** | **1023.31** |
| 1-d | 1412.90* | -23.62* | 59.82* | 18.20* | 982.94* |

## B. Experiment 2 - Mixing or Homogeneous Patterns

Experiment 1 revealed that conservative–stable platoons achieved the best overall performance in terms of delay. In addition, two patterns were observed: i) conservative drivers, owing to their lower wave speeds, tend to generate larger breakdown areas; ii) aggressive drivers produce shorter and less frequent breakdowns, but their higher intensity often results in greater delays and higher fuel consumption.

Experiment 2 therefore examines mixed platoons, first testing six alternating style pairs and then varying the proportions of the most effective pair to assess ratio effects.

*1) General Indicators*

TABLE XIII: Wave speed, standstill time, and delay for each platoon composition (* denotes the difference is significant ($p < 0.001$) compared to the reference level 2-b)

| Exp. | $\lvert w_{\text{dec}}\rvert$ (m/s) | $\lvert w_{\text{acc}}\rvert$ (m/s) | Standstill (s) | Delay (s) | $FC$ (L/km) |
|---|---|---|---|---|---|
| | t-test | | WLS + t-test | | |
| 2-a | +0.20* | +0.28* | +44* | +23* | +0.24* |
| **2-b** | **5.51** | **5.61** | **321.9** | **492.0** | **0.98** |
| 2-b-1 | -0.20* | -0.17* | +17* | -4 | +0.45* |
| 2-b-2 | -0.09 | +0.02 | -3 | -2 | +0.02 |
| 2-b-3 | +0.06 | +0.11 | -19* | +1 | -0.31* |
| 2-c | +0.38* | +0.49* | +6 | +20* | -0.35* |
| 2-d | -0.07 | -0.10 | +71* | +10* | +1.19* |
| 2-e | +0.31* | +0.40* | +81* | +32* | +0.69* |
| 2-f | +0.09 | +0.19* | +40* | +16* | +0.41* |
| 1-a | +0.21* | +0.24* | -39* | +1 | -0.55* |
| 1-b | +0.17 | -0.09 | +129* | +35* | +1.86* |
| 1-c | -0.34* | -0.32* | +29* | -10* | +0.86* |
| 1-d | +0.67* | +0.73* | +43* | +31* | -0.03 |

Table XIII summarizes the results (based on $15 \times 10$ runs). ANOVA confirms significant effects across indicators, with pairwise differences identified through HSD tests.

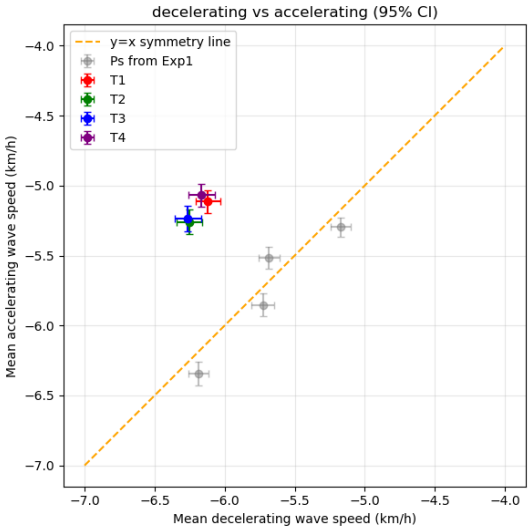


Fig. 15: Wave speed comparison (Margins indicate the 95% CI.)

The results in Table XIII highlight several key findings. First, P3 drivers act as a *stabilizer*: higher proportions consistently reduce delay, though the effect is nonlinear and becomes most pronounced beyond 75%. Second, P2 drivers operate as a *drag*, with P2-dominated platoons (2-d, 2-e) showing the longest delays and standstills. Third, standstill time and delay are not directly coupled, as P3-heavy platoons may experience long standstills yet recover faster overall. Finally, heterogeneous platoons—particularly P1–P3 mixes—perform comparably to the best homogeneous groups but do not substantially outperform them.

Considering fuel use, a trade-off emerges. While P3 stabilizes traffic, its tendency to dwell in low-speed regimes raises $FC$, especially once its share exceeds 50%. By contrast, P1 yields the most favorable $FC$. Mixed P1–P3 platoons thus strike an effective balance: moderate shares of P3 reduce delay without significantly increasing fuel consumption.

### C. Experiment 3 - Impact of Truck in Platoon

Before examining mixed truck–car patterns, it is useful to first analyze homogeneous truck platoons as reference profiles. To this end, three setups are designed, each consisting exclusively of one driving style with a truck as the leader (see Table VI).

The style transformer model accommodates vehicle length as an input, so that agents behave as trucks when assigned a larger length (9 m in this study, denoted by T), in contrast to passenger cars (P).

#### 1) Characteristics of Truck Platoon

Figure 15 summarizes the main findings. Homogeneous truck platoons show a greater gap between decelerating and accelerating wave speeds than passenger-car platoons, meaning that congestion propagates upstream more rapidly but dissipates more slowly. Their accelerating wave speeds are markedly lower, reflecting the slower restart of trucks. By contrast, decelerating wave speeds are slightly higher: despite weaker braking performance, truck drivers' elevated vantage points and professional training allow earlier anticipation and reaction, leading to faster upstream propagation.

TABLE XIV: Difference of indicators compared to reference level (reference group: 3-a, * denotes the difference is significant ($p < 0.01$) compared to the reference level)

| Experiment | $w^{\text{dec}}$ (m/s) | $w^{\text{acc}}$ (m/s) | Standstill (s) | Delay (s) | FC (L/km) |
|---|---|---|---|---|---|
| | t-test | | WLS + t-test | | |
| **3-a** | **6.12** | **5.11** | **21** | **9** | **3.22** |
| 3-b | +0.13* | +0.15* | -38* | +1 | -0.84* |
| 3-c | +0.14* | +0.12 | -41* | -19* | -0.69* |
| 3-d | +0.05 | -0.05 | -7* | -16* | +0.13* |

Meanwhile, other general indicators in Table XIV show two key points:

1) Consistent with earlier findings, a conservative–stable style (T3) proves most effective in reducing overall delay in truck platoons.
2) Fuel consumption is generally higher for trucks than for passenger cars; notably, T3, despite yielding the lowest delay, incurs the highest fuel use due to extended operation in low-speed regimes.

### D. Truck Placement: Clustered or Dispersed

TABLE XV: Impact of truck placement (e,f: clustered, g,h: dispersed)

| Exp. | Composition | $w_{\text{dec}}$ (km/h) | $w_{\text{acc}}$ (km/h) | ST (s) | Delay (s) | FC (L/km) |
|---|---|---|---|---|---|---|
| | | t-test | | WLS + t-test | | |
| 3-e | $[P]^4[T]^8[P]^4$ | **6.0** | **5.4** | **442** | **632** | **4.25** |
| 3-f | $[T]^8[P]^8$ | +0.2 | +0.5* | +21* | +65* | -0.6* |
| 3-g | $[T,P,\cdots]^{16}$ | -0.5* | +0.4 | -35* | +9* | -1.1* |
| 3-h | $\{T(50\%),P\}^{16}$ | -0.4* | +0.3 | -43* | +0 | -1.2* |

This experiment investigates the impact of truck allocation within a platoon. The results (Table XV) yield the following insights:

- **Clustered formations:** Placing trucks together can reduce delay under specific conditions (e.g., 3-e with trucks in the middle), but the benefit is highly sensitive to cluster position (compare 3-e vs. 3-f). In real traffic, where truck positions cannot be controlled, such delay benefits are unreliable, and 3-e further incurs the highest fuel consumption due to prolonged low-speed operation.
- **Clustering greatly increases standstill time**, raising safety concerns as congestion propagates further upstream and lasts longer. In single-lane traffic, clustered trucks trapped among passenger cars may therefore amplify rather than mitigate breakdowns and delays.
- **Dispersed formations:** Random or dispersed allocations yield more resilient dynamics. Although delays may not be as low as in 3-e, shorter standstills indicate milder breakdowns and faster recovery. Alternating formations are less robust, making random dispersion the most

practical and reliable strategy under high-demand, single-lane conditions.

- **Macroscopic effects:** With a fixed truck proportion, no significant differences are observed in fundamental diagram parameters ($a$, $b$, $k_{\mathrm{crit}}$), implying that macroscopic wave dynamics depend on truck share rather than specific allocation.

## VIII. Conclusion

This study presented a data-driven framework for modeling car-following (CF) behavior by embedding driving style as a latent, time-independent factor. A curated dataset enabled robust model training, and a Transformer-based style embedding model jointly learned behavioral traits and CF dynamics. The embeddings aligned with interpretable metrics (e.g., headway, reaction time) and modulated CF behavior systematically, supporting a continuous view of driving style.

Style-aware CF models demonstrated superior accuracy and robustness in long-term predictions compared to IDM, LSTM, and standard Transformers, leveraging style embeddings to reduce overreaction to local fluctuations. This integration of behavioral interpretability and predictive performance offers a novel lens on driver heterogeneity in traffic flow.

Furthermore, the style embedding reveals four distinct groups along two principal axes—aggressive–timid and stable–unstable—providing a foundational taxonomy for the subsequent platoon experiments.

Building on the CF model, platoon-scale simulations explored how style heterogeneity shapes traffic breakdowns. Using both microscopic (delay, wave speed) and macroscopic (fundamental diagram) metrics, the experiments revealed that:

i) Aggressive drivers generate short, fast waves; conservative ones produce longer, slower waves. Stability reduces delays but raises fuel use due to extended low-speed driving.

ii) Stability-fuel efficiency trade-offs emerge. Stabilizing drivers (over 50%) cut delays but increase fuel use. A balanced mix below this threshold yields optimal outcomes.

iii) Trucks amplify deceleration waves, slowing recovery. Dispersed truck placements enhance resilience, while clustering increases fragility, especially under heavy traffic.

This study demonstrates that microscopic driver heterogeneity significantly shapes macroscopic traffic dynamics through three key mechanisms: stability mitigates congestion but increases fuel use; aggressiveness worsens breakdowns despite localized efficiency; and spatial arrangement—particularly truck dispersion—affects system resilience. Methodologically, it offers a reproducible training and testing framework combining a style-aware car-following model. Practically, it recommends promoting stability-oriented driving and avoiding truck clustering to enhance flow efficiency and reduce fuel consumption.

## References


Mark Brackstone and Mike McDonald. Car-following: a historical review. *Transportation Research Part F: Traffic Psychology and Behaviour*, 2(4):181–196, 1999. ISSN 1369-8478. doi: 10.1016/S1369-8478(00)00005-X.

Danjue Chen, Soyoung Ahn, Jorge Laval, and Zuduo Zheng. On the periodicity of traffic oscillations and capacity drop: The role of driver characteristics. *Transportation Research Part B: Methodological*, 59:117–136, January 2014. ISSN 0191-2615. doi: 10.1016/j.trb.2013.11.005.

Kuan-Ting Chen and Huei-Yen Winnie Chen. Driving Style Clustering using Naturalistic Driving Data. *Transportation Research Record*, 2673(6):176–188, June 2019. ISSN 0361-1981. doi: 10.1177/0361198119845360. Publisher: SAGE Publications Inc.

Linsen Chong, Montasir M. Abbas, Alejandra Medina Flintsch, and Bryan Higgs. A rule-based neural network approach to model driver naturalistic behavior in traffic. *Transportation Research Part C: Emerging Technologies*, 32:207–223, July 2013. ISSN 0968-090X. doi: 10.1016/j.trc.2012.09.011.

Chiara Colombaroni and Gaetano Fusco. Artificial Neural Network Models for Car Following: Experimental Analysis and Calibration Issues. *Journal of Intelligent Transportation Systems*, 18(1):5–16, January 2014. ISSN 1547-2450. doi: 10.1080/15472450.2013.801717.

L. C. Edie. Discussion of traffic stream measurements and definitions. pages 139–154, 1965.

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, September 2023. ISSN 0968-090X. doi: 10.1016/j.trc.2023.104272.

Bryan Higgs and Montasir Abbas. Segmentation and clustering of car-following behavior: Recognition of driving patterns. *IEEE Transactions on Intelligent Transportation Systems*, 16(1):81–90, 2015. doi: 10.1109/TITS.2014.2326082.

Sepp Hochreiter and Jurgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL https://doi.org/10.1162/neco.1997.9.8.1735.

Xiuling Huang, Jie Sun, and Jian Sun. 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, 2018. ISSN 0968-090X. doi: 10.1016/j.trc.2018.07.022.

Jorge A. Laval and Ludovic Leclercq. A mechanism to describe the formation and propagation of stop-and-go waves in congested freeway traffic. *Philosophical Transactions of the Royal Society A*, 368(1928):4519–4541, 2010. doi: 10.1098/rsta.2010.0138.

Lijing Ma and Shiru Qu. A sequence to sequence learning based car-following model for multi-step predictions considering reaction delay. *Transportation Research Part C: Emerging Technologies*, 120:102785, November 2020. ISSN 0968-090X. doi: 10.1016/j.trc.2020.102785.

Melanie Mitchell. *An Introduction to Genetic Algorithms*. MIT Press, Cambridge, MA, 1996. ISBN 0-262-13316-4. doi: doi.org/10.7551/mitpress/3927.001.0001.

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, September 2021. ISSN 0968-090X. doi: 10.1016/j.trc.2021.103240.

Marcello Montanino and Vincenzo Punzo. Trajectory data reconstruction and simulation-based validation against macroscopic traffic patterns. *Transportation Research Part B: Methodological*, 80:82–106, 2015. ISSN 0191-2615. doi: 10.1016/j.trb.2015.06.010.

Saskia Ossen and Serge P. Hoogendoorn. Heterogeneity in car-following behavior: Theory and empirics. *Transportation Research Part C: Emerging Technologies*, 19(2):182–195, 2011. ISSN 0968-090X. doi: 10.1016/j.trc.2010.05.006. Emerging theories in traffic and transportation and methods for transportation planning and operations.

Vasileia Papathanasopoulou and Constantinos Antoniou. Towards data-driven car-following models. *Transportation Research Part C: Emerging Technologies*, 55:496–509, 2015. ISSN 0968-090X. doi: 10.1016/j.trc.2015.02.016. Engineering and Applied Sciences Optimization (OPT-i) - Professor Matthew G. Karlaftis Memorial Issue.

Rijkswaterstaat. Rapportage Rijkswegennet. 3e periode 2023, 1 september - 31 december. Technical report, Rijkswaterstaat, 2024. In Dutch.

Fridulv Sagberg, Selpi, Giulio Francesco Bianchi Piccinini, and Johan Engström. A Review of Research on Driving Styles and Road Safety. *Human Factors*, 57(7):1248–1275, November 2015. ISSN 0018-7208. doi: 10.1177/0018720815591313. Publisher: SAGE Publications Inc.

Anshuman Sharma, Zuduo Zheng, and Ashish Bhaskar. A pattern recognition algorithm for assessing trajectory completeness. *Transportation Research Part C: Emerging Technologies*, 96:432–457, November 2018. ISSN 0968-090X. doi: 10.1016/j.trc.2018.09.027. URL https://www.sciencedirect.com/science/article/pii/S0968090X18305904.

Fulvio Simonelli, Gennaro Nicola Bifulco, Valerio De Martinis, and Vincenzo Punzo. Human-like adaptive cruise control systems through a learning machine approach. In Erel Avineri, Mario Köppen, Keshav Dahal, Yos Sunitiyoso, and Rajkumar Roy, editors, *Applications of Soft Computing*, pages 240–249, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg. ISBN 978-3-540-88079-0.

SWOV. Traffic congestion and roadworks. *SWOV-Fact sheet, May 2022*, May 2022. Place: The Hague Publisher: SWOV.

Martin Treiber, Ansgar Hennecke, and Dirk Helbing. Congested traffic states in empirical observations and microscopic simulations. *Phys. Rev. E*, 62:1805–1824, Aug 2000. doi: 10.1103/PhysRevE.62.1805.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, page 6000–6010, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.

Xiao Wang, Rui Jiang, Li Li, Yi-Lun Lin, and Fei-Yue Wang. Long memory is important: A test study on deep-learning based car-following model. *Physica A: Statistical Mechanics and its Applications*, 514:786–795, January 2019. ISSN 0378-4371. doi: 10.1016/j.physa.2018.09.136.

S. K. Zegeye, B. De Schutter, J. Hellendoorn, E. A. Breunesse, and A. Hegyi. Integrated macroscopic traffic flow, emission, and fuel consumption model for control purposes. *Transportation Research Part C: Emerging Technologies*, 31:158–171, 2013. ISSN 0968-090X. doi: https://doi.org/10.1016/j.trc.2013.01.002.

Zuduo Zheng, Soyoung Ahn, Danjue Chen, and Jorge Laval. Freeway traffic oscillations: Microscopic analysis of formations and propagations using wavelet transform. *Transportation Research Part B: Methodological*, 45(9):1378–1388, 2011. ISSN 0191-2615. doi: 10.1016/j.trb.2011.05.012. Select Papers from the 19th ISTTT.

Mofan Zhou, Xiaobo Qu, and Xiaopeng Li. A recurrent neural network based microscopic car following model to predict traffic oscillation. *Transportation Research Part C: Emerging Technologies*, 84:245–264, 2017. ISSN 0968-090X. doi: 10.1016/j.trc.2017.08.027.