



OPEN

Adaptive physics-informed trajectory reconstruction exploiting driver behavior and car dynamics

Michail A. Makridis & Anastasios Kouvelas

As more and more trajectory data become available, their analysis creates unprecedented opportunities for traffic flow investigations. However, observed physical quantities like speed or acceleration are often measured having unrealistic values. Furthermore, observation devices have different hardware and software specifications leading to heterogeneity in noise levels and limiting the efficiency of trajectory reconstruction methods. Typical strategies prune, smooth, or locally modify vehicle trajectories to infer physically plausible quantities. The filtering strength is usually heuristic. Once the physical quantities reach plausible values, additional improvement is impossible without ground truth data. This paper proposes an adaptive physics-informed trajectory reconstruction framework that iteratively detects the optimal filtering magnitude, minimizing local acceleration variance under stable conditions and ensuring compatibility with feasible vehicle acceleration dynamics and common driver behavior characteristics. Assessment is performed using both synthetic and real-world data. Results show a significant reduction in the speed error and invariability of the framework to different data acquisition devices. The last contribution enables the objective comparison between drivers with different sensing equipment.

Technological advancements in road transport sensing technologies, driver assistance systems, and traffic-vehicle monitoring systems generate detailed trajectory data that produce unique insights on different topics, i.e., driver dynamics understanding, emissions and fuel consumption estimations, traffic and vehicle safety investigations, traffic management solutions, and maneuver monitoring. Usually, trajectory reconstruction is performed through pre-processing for outlier detection and further signal filtering. However, reconstructing different signals that might differ in noise, frequency, or sensor quality is challenging using a common strategy. Trajectory data are pruned, smoothed, or locally adjusted using heuristics, aiming for results that reflect physically compatible quantities.

There is a wide range of data collection and sensing technologies such as drones, cameras, Global Positioning Systems (GPS), cellular data, Bluetooth data, and others. In principle, vehicle trajectory monitoring involves regularly observing the vehicle's position at discrete points in time. Then usually, speed and acceleration are derived from the fusion of position measurement with other sensors (e.g., accelerometers, gyroscopes, radars) or direct derivation of positions and speeds, respectively. Different data acquisition devices employ different characteristics (accuracy, noise, sampling frequency), and raw data, in most cases, if not all, need post-processing.

Vehicle trajectory extraction, reconstruction, filtering, or smoothing are closely related and have attracted research interest more than half a century ago with the seminal works of¹ (i.e., Wiener filter), and² (i.e., Rauch-Tung-Striebel smoother). It is an important topic that can enable investigations on various other topics, e.g., emissions or energy demand^{3,4}, traffic dynamics⁵⁻⁸, modeling and control⁹⁻¹¹, control¹¹, connected and automated vehicles and driver understanding^{12,13}, while this list is not exhaustive.

At the same time, more and more datasets with experimental observations have become publicly available, and nowadays, it becomes easier to organize experimental campaigns and study complex phenomena empirically. Filtering techniques are applied to the raw data to remove noise and obtain quality measurements. In the pNEUMA dataset¹⁴, the post-processing is performed by an Advanced Kalman filter without further information on the model or parameters. In the HighHD dataset¹⁵, the RTS algorithm² is applied for post-processing. The OpenACC dataset¹⁶ contains different experiments using either local regression or moving average, while other works^{17,18} use moving average for trajectory filtering.

Recently, Toledo et al.¹⁹ used weighted local regression to derive a smooth time-continuous trajectory function, thus proposing a technique independent of the sampling frequency. Different window sizes and polynomial

Department of Civil, Environmental and Geomatic Engineering, ETH, 8093 Zurich, Switzerland. ✉ email: michail.makridis@ivt.baug.ethz.ch

orders are tested, observing that fluctuations decrease with an increase in the window size, but increase with higher orders of the fitted polynomial. In another interesting work²⁰, the authors have proposed a four-step approach, first removing extreme positional errors, then applying a low-pass filter, third reconstructing the trajectory locally, and finally, removing the residual noise with another low-pass filter. Furthermore, within the broader family of frequency-related techniques, Fard et al.²¹ have proposed a robust wavelet-based two-step method for vehicle trajectory reconstruction. Results have been presented for the well-known NGSIM database, i.e., the Next Generation SIMulation program of the US Department of Transportation in 2002. Additionally, Punzo et al.²² have proposed a Kalman-based filtering technique to reconstruct car-following trajectories and maintain platoon consistency. Finally, working on trajectory extraction from unmanned aerial vehicles, Chet et al.²³ have proposed a novel methodological framework for automatic and accurate vehicle trajectory extraction from aerial videos.

Estimating the inherent noise in experimental data is challenging, although some data acquisition devices provide such analysis. Validation of any applied methodology is impossible without ground-truth information²¹. A common approach is to focus on tackling irregularities, i.e., abnormalities that can be detected with high certainty to comply with comfortable and maximum acceleration standards. For example, an observation of an extreme acceleration (e.g., 10 [m/s²]) or an extreme deceleration during car-following (e.g. -10 [m/s²]) are unarguably unrealistic values. Consequently, anything below (above) that observed acceleration (deceleration) is closer to the real value. The above example can also be applied to movements or speeds, but due to differentiation, such irregularities are more prominent in acceleration. The second indication of abnormal patterns is the high local acceleration variance, even under stable conditions^{19–21}.

Without ground-truth data, the trajectories are filtered to a point where the inferred accelerations, speeds, and movements converge to values that seem reasonable from physics and experience points of view. Usually, there are commonly adopted parameter values per methodology to adjust the filtering strength²⁴. Additionally, Montanino et al.²⁰ validate their approach from a platoon consistency viewpoint, which is also important from a traffic flow perspective.

The vehicle's power dynamics and general driving characteristics directly impact how a vehicle moves. This additional source of information, already exploited in modeling and simulation activities^{25–27}, can be exploited in trajectory reconstruction. The synergistic combination of mathematical or physical models and data^{28–32} is the general idea and motivation behind the design of the proposed framework.

We know from vehicle power dynamics that the maximum vehicle acceleration can be modeled as a function of the vehicle's speed. Consequently, acceleration values that seem perfectly normal at lower speeds are unrealistic at higher speeds³³. Driving behavior (for the longitudinal direction) poses additional constraints in the feasible domain of observed accelerations. Real-world data show that human drivers systematically exploit less than 50% of the vehicle's acceleration capacity at any given speed³⁴. Therefore, we argue that even if the resulting acceleration values seem feasible from a vehicle's power perspective, they might rarely be observed in real-world measurements.

This paper proposes a novel approach that builds on the benefits of traditional techniques and exploits their performance in trajectory reconstruction by injecting additional information about vehicle dynamics and driver behavior (for the longitudinal direction) in a straightforward yet rigorous manner. We propose a threshold-free interactive approach that derives the filtering strength automatically. The main idea is to exploit model-based approaches and reconstruct vehicle trajectories to the point that derived speeds and accelerations not only make sense from a physics point of view but are compliant with the vast majority of the observed vehicles/fleet and drivers on road transport systems.

The framework employs an iterative procedure that includes three main components: (a) the Vehicle Dynamics Constraint process (*VDC*), (b) the Driver Dynamics Compliance process (*DDC*), and (c) the Noise Reduction process (*NR*). The first component models the observed vehicle's acceleration capacity as a speed function. The *VDC* box constrains all accelerations outside this acceleration capacity space. The second component models the acceleration capacity of ordinary observed drivers as a speed function. The vehicle specifications that can be retrieved from online open databases are needed for the first two components, i.e., gearbox, maximum torque, mass, etc., to model the acceleration capacity functions. The acceleration capacity functions for the vehicle's capability and the ordinary behavior of drivers are modeled using the MFC model, considering also observed findings in the literature^{25,34}. For implementation details, we refer the reader to the corresponding paper and the publicly available library [<https://pypi.org/project/co2mpas-driver/>]. The third component filters the signal with a predetermined filtering method. The filtering strength (also reported in this work as magnitude) is automatically detected through an iterative process that monitors and minimizes the local acceleration variance²⁹ and is described in detail in the methodological section. It should be noted that the entire adaptive process is automatic and invariant to the quality of the observed data.

The framework's assessment is performed by utilizing synthetic and real-world data. Highly accurate differential GPS observations are used as reference trajectories, i.e., ground truth. Five levels of Gaussian noise are added, leading to 5 noisy trajectory datasets. Then, the proposed approach is tested on its ability to reconstruct the reference trajectories. Additionally, two mobile devices obtain real-world data on the same vehicle trajectory. The two devices have very different specifications resulting in visually different observations. The proposed methodology is applied to each signal separately and reduces the signal differences to the extent that the framework can be considered invariant to the acquisition device. Therefore, the results of the real-world campaign are promising for comparative investigations between different drivers. Figure 1 illustrates a high-level summary of the proposed framework.

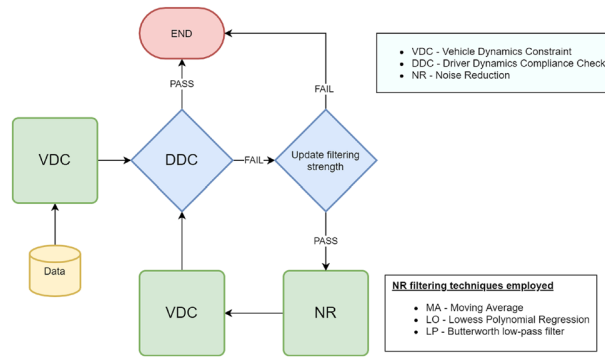


Figure 1. Flowchart for the proposed physics-informed framework.

Results

Synthetic and real-world data are used to assess the robustness of the solution. Synthetic noisy data are generated by adding noise to observations obtained using highly accurate differential GPS devices (the AstaZero campaign of the OpenACC dataset¹⁶). Data from this campaign are used as a reference, i.e., ground truth. Five noisy synthetic trajectories are generated, assuming the reference data are noise-free. These trajectories, namely $N1, N2, N3, N4, N5$, have zero mean Gaussian noise and increasing standard deviation from 0.05 to 0.25 [m/s] at frequency of 10 [Hz]. The noise levels reflect the error estimates given by commercial data acquisition systems, such as U-blox. Three state-of-the-art filtering techniques are tested within the framework in the *NR* component (see Fig. 1), i.e., Moving Average (*MA*), Lowess Polynomial Regression algorithm (*LO*), and Butterworth filter (*LP*), while other similar filtering approaches can also be hosted. We compare the performance of each method when used independently with parameters suggested in the literature versus the same method when employed within the proposed framework, where the filtering strength is adaptively adjusted. When the methods are used independently, fixed parameters are assumed, a time window of 3s for *MA*, a 7-step window for *LO*, and a cut-off frequency of 0.75 [Hz] for *LP* (see²⁰ for details). When the above techniques are used inside the proposed framework, their notation is *pMA*, *pLO*, and *pLP*, respectively, and their filtering magnitudes (window, steps, or cut-off frequencies) are inferred adaptively. Additional insights that demonstrate the robustness of the proposed methodology in comparison to more sophisticated methodologies for trajectory reconstruction (here used the Fard et al. technique²¹) are given in the Appendix document.

Figure 2a,b illustrate the I_{RMSE} and I_{MAE} indicators for each scenario, that is, the root mean square error and mean absolute error of speed indices, as they are described in the methodology. The incorporation of each one of the three tested methods (*MA*, *LP*, *LO*) within the proposed framework (*pMA*, *pLP*, *pLO*) leads to a significantly improved reconstruction. Higher noise levels correspond to higher errors in comparison to the ground truth. The *LO* and *pLO* methods are the best performing and independent of the data frequency, which is a considerable advantage¹⁹. The *LO* performs well even with fixed window size, and the efficiency is comparable to the more complex counterparts of *pMA* and *pLO*. Within the proposed framework (*pLP*), the error decreases drastically, and it outperforms all the other methods compared to the other methods. The *MA* and *pMA* are fast, overall reliable, and almost independent of the noise levels, while *LP* and *pLP* are not recommended for this task.

The proposed framework considers the specifications of the vehicle known. However, this is not always possible in real-world campaigns. Consequently, for large-scale application of the proposed framework, average vehicle dynamics²⁵ using representative vehicles from Euro Car Segments, as they are defined by European

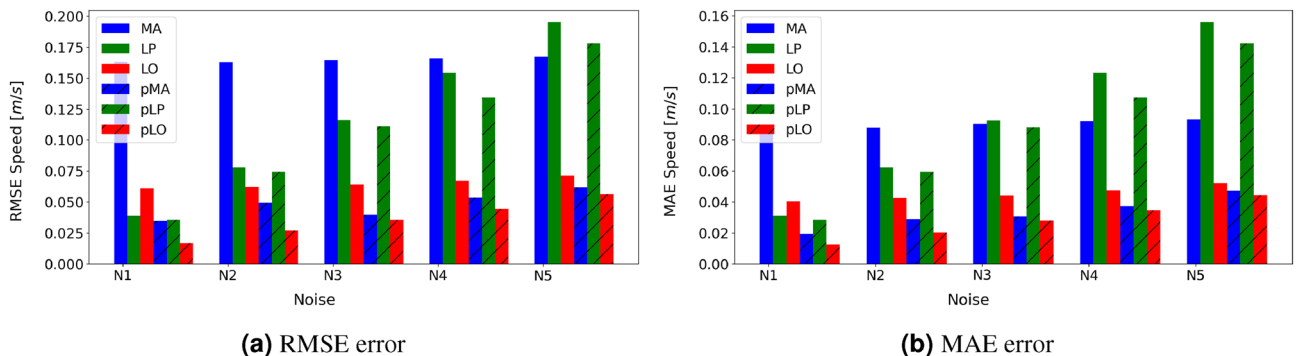


Figure 2. The RMSE and MAE errors for the three reconstruction methods with *pMA*, *pLP*, *pLO* and without (*MA*, *LP*, *LO*) the proposed framework.

Commission policy³⁵ can be employed. Alternatively, a more generic vehicle dynamics model²⁶ with average vehicle specifications can also substitute the proposed MFC.

Table 1 shows the inferred optimal window size and cut-off frequency as computed adaptively by the proposed framework. The *MA* has been applied with fixed window, i.e., 3 [s], the *LO* with fixed steps, i.e., 7 (or 0.7 [s] at 10 [Hz] data), and *LP* with fixed frequency (0.75 [Hz]), which are commonly suggested values in the literature.

The application of the proposed adaptive approach improves the efficiency of each method with ordinary filtering strength by up to 80% in terms of error reduction. However, aggregated indicators do not give insights about the resulting microscopic dynamics after the trajectory reconstruction, essential for accurate traffic flow and energy-related conclusions⁴.

Insights on the reconstructed microscopic dynamics follow in Fig. 3. Figure 3a,c,e,g,i illustrate the acceleration over speed values for each synthetic dataset *N1*, *N2*, *N3*, *N4*, *N5* respectively (red dots). The blue dots correspond to the filtered acceleration speed values after the application of the *LO* method. Figure 3b,d,f,h,j show results with the proposed framework. The red dots correspond to the synthetic datasets, the green dots to the corrections applied by the *VDC* component, and the black dots are the output of the proposed framework with the *pLO* method. Figure 3k shows the reference data, while Fig. 3l provides the legend for all sub-figures. The comparison of Fig. 3a,j with Fig. 3k demonstrates some interesting findings regarding the ability of each method to capture the observed acceleration dynamics.

An oscillation of acceleration values for speeds between 0 and 15 [m/s] is prominent in the ground truth data. This oscillation is mostly filtered by the individual *LO* application, while it is still observable in the results of the proposed physics-informed framework, i.e., *pLO*. Furthermore, the main body of acceleration observations between speeds 15 [m/s] and 25 [m/s] for the proposed technique is visually close to the ground truth data. The *LO* method homogenizes this area, resulting in smoother but distorted dynamics compared to ground truth data. The results show the efficiency of the proposed framework, *pLO* even for high noise levels as in scenarios *N3*–*N5*.

The second campaign took place in Greece. A vehicle's trajectory is observed with two smartphones (*S1* and *S2*), differing significantly in price, operating systems, and specifications. *S1* costs half the price of *S2*, has Android, while *S2* iOS, and records with an average rate of 1s (min/max: 0.04 [s]/2.89 [s]), while *S2* records with the same frequency on average, i.e. 1s but different variance (min/max: 0.95 [s], 2.07 [s]).

Both signals are re-sampled to 1 [Hz] using linear interpolation. A lack of synchronization is most probably due to the internal clock of each device. Cross-correlation is used to estimate the time lag between the signals⁵.

Figure 4 shows the speed and acceleration profiles per measurement device before and after the application of *pLO* method, which was the most efficient according to the results of the first campaign. The observed profiles are shifted in time, as mentioned above, to become synchronized. Both smartphones record the same vehicle trajectory (being with the driver), and the observed speed profiles are shown in the sub-figure Fig. 4a. Then the *pLO* is applied to both speed profiles, and the result is illustrated in the sub-figure Fig. 4b. Visually inspection reveals that most of the differences between the two signals have been smoothed, and it is obvious that they refer to the same measured trajectory. Going one order up, the acceleration observations are shown in the sub-figure Fig. 4c. Here, the differences between the observations of the two devices are more prominent. After the application of the proposed framework (*pLO*), the generated acceleration signals become very similar, pointing to the same measurements; see the sub-figure Fig. 4d. This result is very desirable in trajectory reconstruction. It can provide reliable results for several applications involving user comparisons concerning driving behaviors, fuel consumption profiles, and others.

Table 2 provides a quantitative evaluation for the second campaign, comparing the mean absolute error (MAE) and the median absolute deviation (MAD) between the two observed signals before and after the application of the proposed methodology. The comparative results refer to speed and acceleration values. In terms of speed, the improvement by the adaptive methodology is around 22.7% and 18.2% for MAE and MAD. The corresponding values for acceleration are even higher, i.e., 69.7% and 72.5% for MAE and MAD.

Discussion

The vehicular and driving behaviors are known in traffic engineering as perhaps the most significant factors for raising complexity and leading to the appearance of non-linear phenomena in road transport systems. Until recently, experimental observations regarding detailed vehicle trajectory data have been scarce. During the last decade, technological advances and cost reductions in sensors enabled the organization of large and complex experimental campaigns. Such datasets, many of which are publicly available, provide invaluable insights into traffic engineering topics, provided they have a low error in observations.

	N1	N2	N3	N4	N5
<i>pMA</i>	1.2 [s]	1.5 [s]	0.8 [s]	1.3 [s]	1.1 [s]
<i>pLO</i>	25	31	29	31	37
<i>pLP</i>	0.65 [Hz]	0.7 [Hz]	0.7 [Hz]	0.60 [Hz]	0.65 [Hz]

Table 1. Optimal filtering strength for each technique with the proposed framework. For *pMA* and *pLO* methods, the first number corresponds to the number of observed points, while the number in parenthesis corresponds to its equivalent in terms of time.

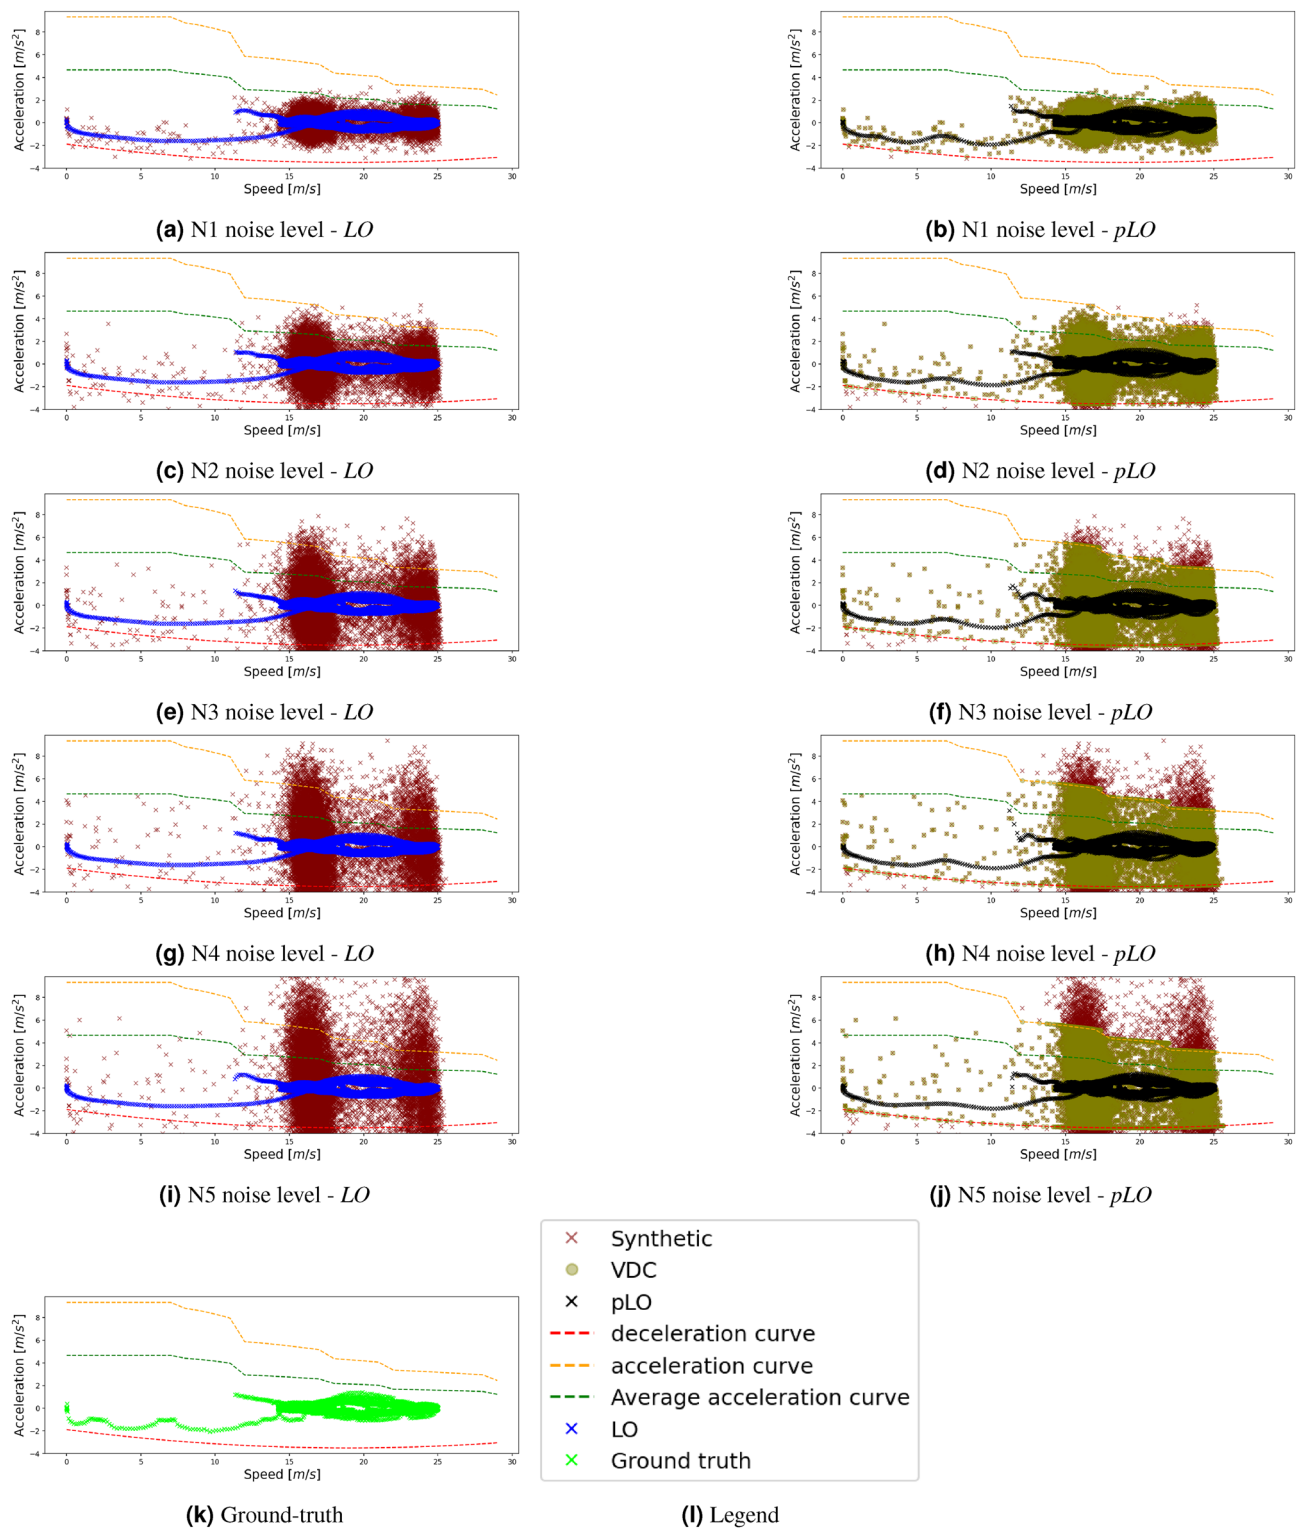


Figure 3. Campaign 1: Comparative results between the Lowess Polynomial Regression with and without the proposed framework for different noise levels.

There is no standardized protocol for designing and executing such experiments. Therefore, each experiment is performed with a different acquisition device, e.g., smartphones, U-blox, OXTS, etc. Such devices have various hardware and software specifications; therefore, the output signals are different even under the same conditions. Noise removal in the measurements is often performed with arbitrarily parametrized filtering approaches, leading to questionable results since no ground-truth reference is available. The main idea is to remove by thresholding obvious outliers, i.e., values with no physical meaning, and then smooth the observation series

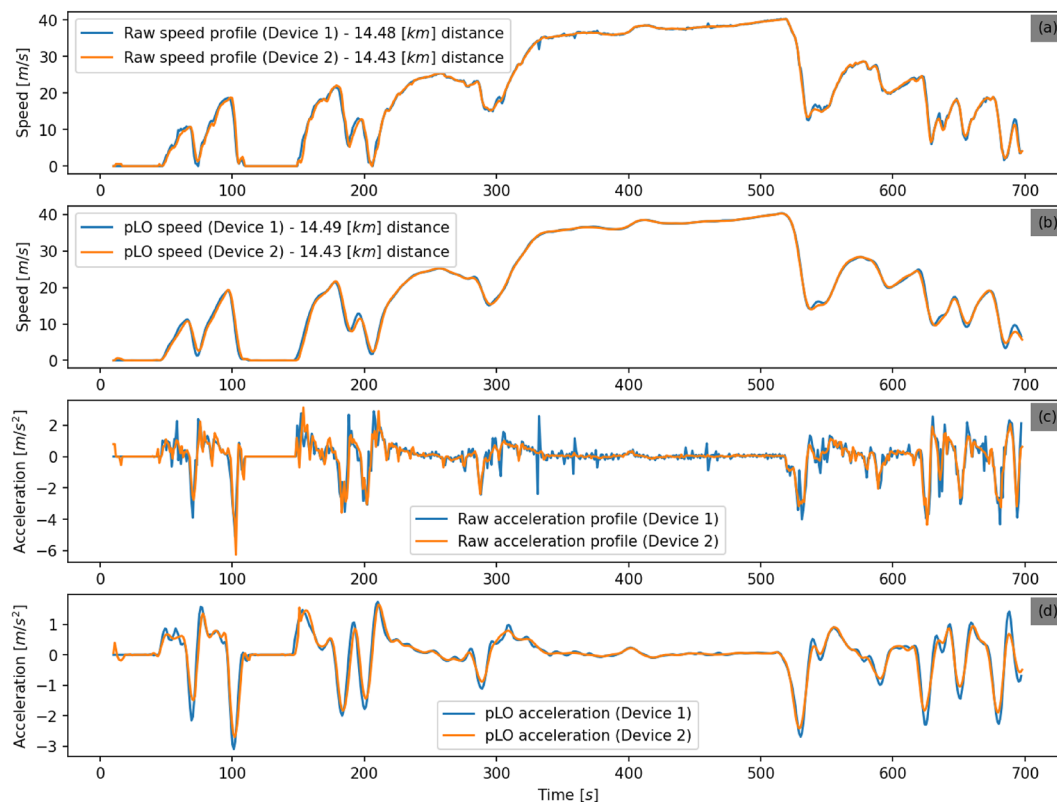


Figure 4. Campaign 2: The speed and acceleration profiles for the same trajectory with two devices, before and after the application of *pLO* method.

	MAE before	MAE after	MAD before	MAD after
Speed [m/s]	0.44	0.34	0.22	0.18
Acceleration [m/s ²]	0.33	0.1	0.16	0.044

Table 2. Mean absolute error and median absolute deviation on speed and acceleration between the two observed signals before and after the application of *pLO* method.

globally or locally (on time or frequency domain) to derive a set of plausible measurements. The above strategy is problematic by design since vehicular and driver dynamics are nonlinear and observations with different devices need custom parametrization that is not obvious.

Exploiting the recent modeling advances in longitudinal vehicle dynamics simulation and driver behavior, the current methodology explicitly considers these two dimensions for trajectory reconstruction. The framework design is flexible and model-agnostic. A convergence process initiates, and an iterative process automatically adjusts the filtering strength based on vehicle and driver dynamic constraints.

Results on synthetic data generated based on low-error reference observations show that the proposed framework remarkably removes outliers and noise, maximizing the efficiency of the employed filtering strategy by automatically parametrizing the filtering strength. Furthermore, real-world observations on the same trajectory with two different devices show significant differences between measurements. The proposed framework is applied in both measured trajectories that visually demonstrate remarkable resemblance afterward, which is normal as they refer to the same experiment.

The above results are unique among trajectory reconstruction methodologies. Compliance with realistic nonlinear vehicle and driver dynamics ensures a fair comparison of observed trajectories and behaviors in various topics such as traffic flow, safety, driver identification, etc. Furthermore, the possibility to identify the same experiment from observations with different noise levels facilitates cross-driver comparisons providing reliable insights into topics such as energy consumption, driver aggressiveness, etc.

The main assumption of this framework is that the vehicle specifications, i.e., gearbox, gear ratio, mass, maximum torque, etc., are considered known for each vehicle trajectory²⁵. This assumption is quite heavy because this information is not always given, especially in large complex experimental campaigns. However, it is shown that the vehicle dynamics can be efficiently clustered based on some main specifications³⁵. The vehicles in the same

cluster demonstrate similar dynamics, so no vehicle-specific details are necessary. Alternatively, there are more generic models than MFC, see for example²⁶, that use average vehicle specifications and are much more flexible to be used within the proposed framework. It is worth noting that the dynamics of electric powertrains differ significantly from vehicles with Internal Combustion Engines, as the first can offer much higher torque from very low speeds. Therefore a different modeling approach should be used for electric vehicles. An implementation of the MFC for electric powertrains is also publicly available³⁶. Finally, an extension or adoption of the proposed framework to capture both lateral and longitudinal dynamics would be interesting as a future work³⁷.

Methods

A physics-informed adaptive framework for trajectory reconstruction is proposed in this work. The algorithm employs three main components, as illustrated in Fig. 1: (a) the Vehicle Dynamics Constraint (VDC) process, (b) Driver Dynamics Compliance (DDC) process, and (c) Noise Reduction (NR) process. The next subsections describe these components individually, while the last subsection here presents the experimental campaigns and scenarios used for assessing and validating the approach, as well as the performance indicators.

Before discussing the individual components, it is important to elaborate on how vehicle and driver dynamics are modeled and considered in this approach. Figure 5 aims to clarify vehicle and driver dynamics from a modeling perspective in an illustrative way. Using a vehicle's specifications (powertrain, engine power, mass, etc.), MFC microsimulation model²⁵, used in this framework, describes the vehicle's continuous acceleration capacity function, namely $a_p(v)$. This function returns the maximum possible acceleration for this specific vehicle at any given speed v . Figure 5 depicts this with the orange line.

Similarly, we compute the continuous minimum comfortable deceleration curve³⁶, namely $d_p(v)$ (non-safety-critical situations). This function provides the minimum possible acceleration for this specific vehicle and a given speed v . Figure 5 depicts this with the red line.

These two curves inscribe our vehicle power domain. In other words, assuming that the employed vehicle dynamics model is accurate, any real acceleration value for a given speed that falls outside this domain is considered unrealistic (infeasible). Although the results in this work refer to specific vehicle models, it should be noted that the proposed methodology can be expanded for vehicle classes, according to Euro Car Segments, defined by the European Commission³⁸, without significant loss of precision.

Furthermore, driver behavior observations reveal that most drivers do not even approach the acceleration capacity of their vehicle for any given speed. Recently, Makridis et al.³⁴ conducted an unbiased experiment with 20 individuals driving freely (without any given instructions) across Europe the same vehicle over one year. The authors proposed a methodology to characterize the aggressiveness of these drivers and possibly cluster them in groups. One of the conclusions in that work has been that in all observations, maximum observed acceleration never exceeded 50% of the vehicle's acceleration capacity for that given speed. Based on this conclusion, we compute the continuous acceleration capacity function of the driver, namely $a_d(v)$ as follows:

$$a_d(v) = d \cdot a_p(v) \quad (1)$$

where v is a given speed and d is a scaling parameter. As mentioned above, in the current work, this parameter is fixed to 0.5, which is the recommended value for common-purpose car-following/driving experiments. In

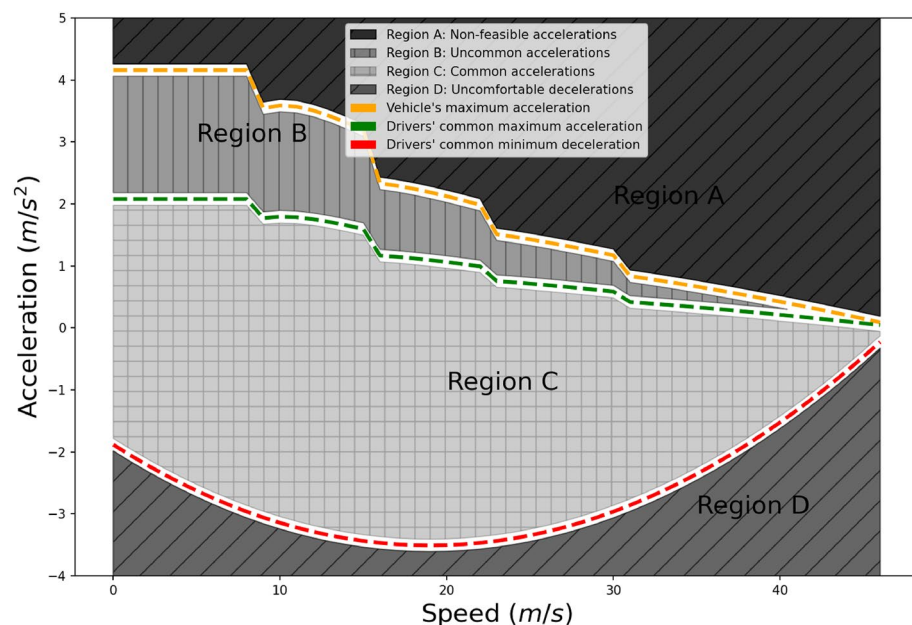


Figure 5. The acceleration-speed domain for a given vehicle and an average driver as computed in^{25,36}.

particular cases, such as experiments with free-flow accelerations to racing events, d should be tuned to values much closer to 1. This function provides the maximum estimated acceleration for most drivers and a given speed v . Figure 5 depicts this threshold with the green line.

Using the above three functions, $a_p(v)$, $d_p(v)$ and $a_d(v)$, we partition the acceleration over the speed domain for all possible observations in four main regions (Fig. 5) as follows:

- Region A: This set, R_A , includes all possible acceleration values for a given speed that the vehicle can not exploit. Any observation inside Region A is considered an outlier and should be box-constrained to a realistic value.
- Region B: This set, R_B , includes all possible acceleration values for a given speed that the vehicle can exploit but are considered not highly probable for common (ordinary) driving conditions. Any observation inside Region B is considered a noisy measurement and therefore indicates the necessity for noise reduction.
- Region C: This set, R_C , includes all possible acceleration values for a given speed that the vehicle can exploit and can potentially correspond to ordinary driving conditions. Any observation that lies inside Region C is acceptable as an accurate measurement. Therefore, there is no way to assess it and evaluate if the measurement corresponds to the actual value or not.
- Region D: This set, R_D , includes all possible acceleration values for a given speed outside a driver's comfortable deceleration values. Any observation inside Region D is considered an outlier and should be box-constrained to a realistic value.

In the context of the current application of reconstructing vehicle trajectories, we denote $x(t)$ as the time series of position measurements for a given vehicle. Moreover, s is the time-step, and $f = 1/s$ is the sampling frequency. Correspondingly, speed and acceleration measurements are obtained with derivation, namely $\dot{x}(t)$ and $\ddot{x}(t)$. Additionally, if N_A, N_B, N_C, N_D are the observations that fall in the corresponding regions R_A, R_B, R_C and R_D , respectively, then, $N_A + N_B + N_C + N_D = N$. Finally, the present study can be employed for electric vehicles that demonstrate different power characteristics by utilizing the corresponding version of the MFC model³⁶.

Vehicle dynamics constraint. The goal of VDC is to ensure that all the observations in regions R_A and R_B are box-constrained to physically possible values, i.e., orange and red lines at the boundaries of regions R_B and R_C in Fig. 5. An advantage of the proposed framework over existing works is that it constrains the outliers nonlinearly. It is a common phenomenon in experiments to produce outliers, values that do not make any sense from the point of view of physics. Such values may appear due to different factors, such as weather conditions, road geometry, sensor errors, malfunctions, and others. Hard thresholding is commonly used in the literature toward this scope, i.e., removing accelerations above 5 [m/s²] and below -8 [m/s²]. However, the application of a horizontal threshold is not efficient. For instance, an acceleration of 3 [m/s²] can be achieved by a vehicle under low speeds, but it is unrealistic for high speeds, i.e., over 100 [km/h]. Reliable data acquisition systems might capture only a few such observations, i.e., that lie in R_A or R_D , as they could also have their built-in filtering techniques. However, when data are collected with low-accuracy devices (e.g., mobile phones), this process plays an essential role.

The resulted observations set that is constrained by vehicle dynamics is called \ddot{x}_{vdc} and can be described as follows:

$$\ddot{x}_{\text{vdc}}(t) = \begin{cases} a_p(\dot{x}(t)), & \text{if } \dot{x}(t) \in R_A \\ d_p(\dot{x}(t)), & \text{if } \dot{x}(t) \in R_D \\ \ddot{x}(t), & \text{otherwise} \end{cases} \quad (2)$$

This process is applied iteratively inside the proposed methodology for increasing filtering strength, generating a new set of values at each iteration. Therefore, the paper uses an indicator l to describe the loop count wherever necessary.

Driver dynamics compliance. The goal of DDC is to ensure that all observations respect the average driver behavior and automatically assess the filtering magnitude that will be applied by the noise reduction technique. Ultimately, the framework should work independently of the noise reduction methodology and its parameters related to the filtering strength. One of the main problems in traditional trajectory reconstruction techniques is the definition of granularity in the noise reduction processes. For example, in low-pass filtering, the cut-off frequency can vary depending on the noise levels in the raw data and/or sampling frequency. Similarly, in polynomial regression or moving average, the span of the window that will be used for smoothing can heavily distort the observed profile. The proposed compliance check focuses on ensuring the following requirements:

- (RQ1) All acceleration values correspond to common driver aggressiveness levels.
- (RQ2) The variance of sequential accelerations locally in time is close to the variance typically observed in high-accuracy data acquisition systems (e.g., Differential GPS).

The first requirement, RQ1, aims at ensuring that measurements correspond to typical drivers. Typical drivers do not exploit the full power and, thus, acceleration capacity of their vehicle³⁴. The term aggressiveness is used within this work to characterize drivers that accelerate sharper than others, i.e., approaching the vehicle's maximum acceleration for a given speed. According to Fig. 5, all observations should lie in the area below the

orange line, i.e., the maximum possible acceleration of the vehicle at a given speed. In practice, most observations fall around the middle of the area inscribed by the orange and red lines (comfortable deceleration as a function of speed). Observations near the orange line or below the red line are rarely observed, indicating excessive aggressiveness in driving. In our opinion, this is something to be considered in trajectory analysis. Of course, it should be noted that for particular types of experiments (e.g., maximum acceleration from 0 to 100 km/h), this requirement should be relaxed, i.e., the green line in Fig. 5 should be defined closer to the orange one (R_B area becomes smaller). This process validates the ratio $r_l = N_{C_l}/N$, which is the number of observations that fall inside R_C , for iteration l , over the total number of observations N . The ratio r_l is parametrized for every iteration l because the number of values that fall in R_C can differ in every iteration. According to this criterion, an acceptable dataset should contain only a few values outside region R_C . Therefore, this ratio needs to be lower than a fixed threshold, namely d_{th} (here set to 0.05, or 5% of the observations).

The second requirement, *RQ2*, aims to monitor the acceleration signal's local variations. Under real-world non-critical driving conditions, the local variance of accelerations should be relatively low. In the frequency domain, the above irregular pattern is often encountered in empirical observations due to noise, and therefore, low-pass filtered techniques are commonly used to mitigate this effect. Consequently, even if all observations are compatible with vehicle and driver dynamics, the magnitude of local acceleration variations should be low before the application of a smoothing technique and therefore plays a decisive role in the parametrization of the filter's strength (i.e., window or cut-off frequency).

On the other hand, the amount of noise in a signal directly impacts the observed variations. Furthermore, each technique has different efficiency in alleviating these variations, which is not known beforehand. The proposed work estimates the parameters that impact the filter strength based on the above. Assuming that we have N acceleration observations $\ddot{x}(t)$, their local standard deviation $\sigma_{\ddot{x},l}(k)$, for iteration l around k , and a fixed window w_{var} (set to 1.5 s) is computed as follows:

$$\sigma_{\ddot{x},l}(k) = \sqrt{\frac{1}{(2 \cdot w_{var} + 1)} \sum_{i=k-w_{var}}^{k+w_{var}} (\ddot{x}(i) - \mu_k)^2}, \quad (3)$$

where

$$\mu_k = \sqrt{\frac{1}{(2 \cdot w_{var} + 1)} \sum_{i=k-w_{var}}^{k+w_{var}} \ddot{x}(i)}. \quad (4)$$

Since the acceleration values change, the local standard deviation is parametrized per iteration, l . Furthermore, we consider as an indicator the median value of all observed local standard deviations, $\text{med}(\{\sigma_{\ddot{x},l}(k)\}, \forall k)$, namely $\text{med}(\sigma_{\ddot{x},l})$. If we consider that most of the time, the driver (either human or automated) aims at maintaining a constant speed, we can assume that $\text{med}(\sigma_{\ddot{x},l})$ will correspond to a prevalent value. However, standard noise directly impacts individual $\text{med}(\sigma_{\ddot{x},l})(k)$ values and consequently the global $\text{med}(\sigma_{\ddot{x},l})$ indicator. Furthermore, we define the following function:

$$f(w_l) = \frac{(\text{med}(\sigma_{\ddot{x},l-1}) - \text{med}(\sigma_{\ddot{x},l}))}{\text{med}(\sigma_{\ddot{x},l})}. \quad (5)$$

The idea behind function $f(w_l)$ is to automatically determine the filtering magnitude for the noise reduction component based on the reduction of the acceleration's local variability. Specifically, increasing filtering strength at each iteration l , i.e., increasing the window/step size or lowering the cut-off frequency, σ is expected to decrease respectively. At the same time, the rate of $\text{med}(\sigma_{\ddot{x},l})$, which in our discrete system is the normalized difference between $\text{med}(\sigma_{\ddot{x},l-1})$ and $\text{med}(\sigma_{\ddot{x},l})$, decreases as well, towards an optimal result corresponding to the applied technique. However, after passing the optimal threshold, we expect the function f to stop being monotonically decreasing and most likely fluctuate as the filter strength increases further (larger windows or lower cut-off frequencies).

An example of f function trajectory over iterations is shown in Fig. 6a. The trajectory's three markers (red, yellow, and blue) correspond to the optimal, typical, and over-filtered window sizes. The idea is that until the red marker, we have high certainty that the noise decreases as the local standard deviation rate decreases. After that iteration, there is little knowledge about real noise reduction. Typical thresholds, depending on the dataset, might lead to a reasonable noise reduction level as was shown in Fig. 2a,b, but the final result depends on the dataset specifications (noise, frequency, etc.). The proposed technique ensures a satisfactory result and automatic inference of the filter's magnitude. Figure 6b illustrates the acceleration over speed diagrams for the typical, optimal, and over-filtered points, showcasing the robustness of the proposed process.

Concretely, we set the filtering range threshold to the maximum filter range that can ensure that function f remains monotonically decreasing, that is, the size of the window or cut-off frequency w_l for which $f(w_l) > f(w_{l-1})$. To this end, we have:

$$\text{RQ2} \begin{cases} \text{PASS,} & \text{if } r_l \leq d_{th} \wedge f(w_l) \geq f(w_{l-1}) \\ \text{FAIL,} & \text{otherwise} \end{cases} \quad (6)$$

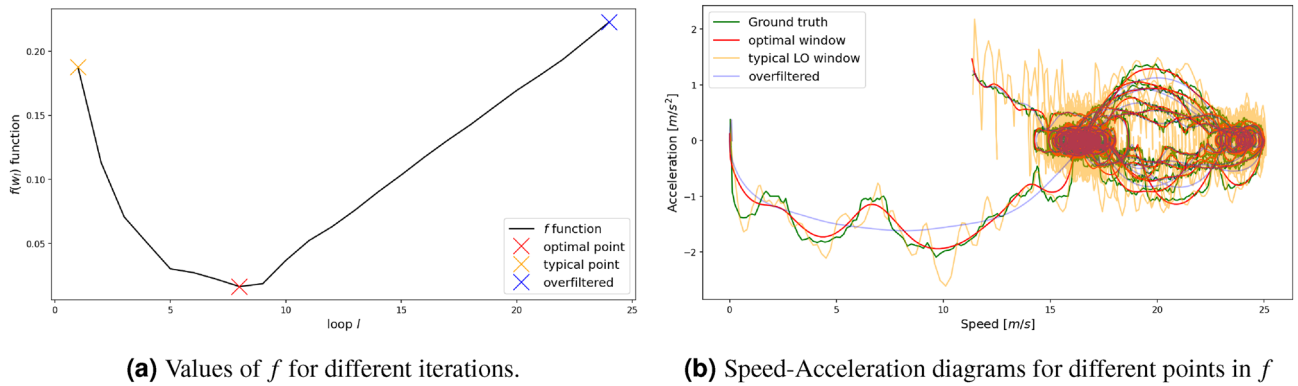


Figure 6. Example of f function over the number of iterations for the LO method with.

Noise reduction and trajectory reconstruction. Three well-known techniques have been employed for trajectory reconstruction and the assessment of the present methodology: a simple Moving Average (*MA*); Lowess Polynomial Regression (*LO*) (window parameter taken from¹⁹); and Butterworth (*LP*) (cut-off parameter taken from²⁰). The window w_{MA} for *MA* looks locally up to 15 observations, $w_{MA} \in \{1, 2, \dots, 15\}$. For *LO* the minimum window size w_{LO} should be sufficient to adequately estimate the derived speeds and accelerations (see also¹⁹), thus $w_{LO} \in \{3, \dots, 15\}$. The order was set to 6. Finally, a first-order Butterworth filter was implemented, and its cut-off frequency f_{LP} decreases progressively down to 0.05 [Hz], $f_{LP} \in \{0.9, \dots, 0.05\}$. We consider that the magnitude of a filtering technique increases as the output is smoother. Therefore, the maximum strength for *MA* and *LO* is with window size 15, while for *LP* with 0.05 cut-off frequency. Table 3 describes the proposed algorithm at a high level. The raw data are pre-processed with *VDC*, and then the algorithm controls if the specifications of *DDC* are met. If not, *NR* is applied, followed by the next iteration. This workflow applies iteratively with increasing noise reduction strength until either *DDC* conditions are met or the maximum *NR* strength is reached.

Performance assessment. The assessment of trajectory reconstruction methodologies is challenging when there is no ground truth data. In this paper, we employ two experimental campaigns with different campaigns to perform our assessment.

Campaign 1: The *AstaZero* experiments described in OpenACC dataset¹⁶ include low-noise position observations from multiple vehicles inside the *AstaZero* test track in Sweden. The measurements have low noise due to the differential GPS data acquisition equipment. This work uses around 25km of a vehicle trajectory in this test track. Because of the initial low noise levels, we consider this reference trajectory as ground truth. Gaussian noises of zero means and increasing standard deviation values are added to create noisy synthetic speed profiles. More specifically, 5 levels of standard deviation in [m/s] are considered in the results, {0.05, 0.1, 0.15, 0.2, 0.25}. The proposed framework is tested for the above three noise reduction methods and five noise levels. Moreover, comparisons are presented for the three noise reduction methods with parameters proposed in the literature.

For comparison indicators, the mean absolute error and root mean squared error on the speed profiles are used;

$$I_{MAE} = \frac{1}{N} \sum_{i=1}^N |\dot{x} - \hat{x}| \tag{7}$$

Initialization
Step 1: Load raw data
Step 2: Initialize filter strengths w_{MA} , w_{LO} or f_{LP}
Step 3: Apply <i>VDC</i>
Step 4: Check <i>DDC</i> : If PASS GoTo Step 9; Else GoTo Step 5
Step 5: Set $w_{MA} = w_{MA} + 1$, $w_{LO} = w_{LO} + 1$ or $f_{LP} = cf_{LP} - 0.05$
Step 6: Apply Noise Reduction (<i>MA</i> , <i>LO</i> or <i>LP</i>)
Step 7: Apply <i>VDC</i>
Step 8: Check <i>DDC</i> : If PASS or filter strength is maximum GoTo Step 9; Else GoTo Step 5
Step 9: END

Table 3. High-level algorithmic workflow.

$$I_{\text{RMSE}} = \sqrt{\frac{1}{N} \sum_{i=1}^N (\dot{x} - \hat{\dot{x}})^2} \quad (8)$$

where N is the number of observations, \dot{x} the ground-truth and $\hat{\dot{x}}$ the reconstructed trajectory.

Campaign 2: This campaign includes the observed trajectories recorded by the sensors of two smartphones via the Phyphox application³⁹. Two different smartphone devices, one less and another more expensive, with Android and iOS operating systems, respectively, have been used during the same experiment. The idea is to assess the consistency of the proposed framework for different sensors and error levels on the same trajectory. We compute the mean absolute error and median absolute deviation between the two signals to provide quantitative assessment values.

Received: 21 November 2021; Accepted: 13 January 2023

Published online: 20 January 2023

References

- Wiener, N. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications* (MIT Press, 1949).
- Rauch, H. E., Tung, F. & Striebel, C. T. Maximum likelihood estimates of linear dynamic systems. *AIAA J.* **3**, 1445–1450. <https://doi.org/10.2514/3.3166> (1965).
- He, Y. *et al.* The energy impact of adaptive cruise control in real-world highway multiple-car-following scenarios. *Eur. Transp. Res. Rev.* **12**, 17. <https://doi.org/10.1186/s12544-020-00406-w> (2020).
- Suarez, J. *et al.* Benchmarking the driver acceleration impact on vehicle energy consumption and CO₂ emissions. *Transp. Res. Part D Transp. Environ.* **107**, 103282. <https://doi.org/10.1016/j.trd.2022.103282> (2022).
- Makridis, M. *et al.* Empirical study on the properties of adaptive cruise control systems and their impact on traffic flow and string stability. *Transp. Res. Rec.* **2674**, 471–484. <https://doi.org/10.1177/0361198120911047> (2020).
- Mattas, K. *et al.* Fuzzy Surrogate Safety Metrics for real-time assessment of rear-end collision risk. A study based on empirical observations. *Accident Anal. Prev.* **148**, 105794. <https://doi.org/10.1016/j.aap.2020.105794> (2020).
- Genser, A., Spielhofer, R., Nitsche, P. & Kouvelas, A. *Critical Ride Comfort Detection for Automated Vehicles*. SVT Working Papers. <https://doi.org/10.3929/ethz-b-000488820> (2021). Accepted: 2021-06-15T10:37:47Z Publisher: IVT, ETH Zurich.
- Chen, D., Laval, J. A., Ahn, S. & Zheng, Z. Microscopic traffic hysteresis in traffic oscillations: A behavioral perspective. *Transp. Res. Part B Methodol.* **46**, 1440–1453. <https://doi.org/10.1016/j.trb.2012.07.002> (2012).
- Kurtc, V. The HighD Dataset: Is this dataset suitable for calibration of vehicular traffic models? In *Traffic and Granular Flow 2019, Springer Proceedings in Physics* (eds Zuriguel, I. *et al.*) 523–529 (Springer, 2020). https://doi.org/10.1007/978-3-030-55973-1_64.
- Sharma, A., Zheng, Z. & Bhaskar, A. Is more always better? The impact of vehicular trajectory completeness on car-following model calibration and validation. *Transp. Res. Part B Methodol.* **120**, 49–75. <https://doi.org/10.1016/j.trb.2018.12.016> (2019).
- Chavoshi, K. & Kouvelas, A. Nonlinear model predictive control for coordinated traffic flow management in highway systems. In *2020 European Control Conference (ECC)* 428–433. <https://doi.org/10.23919/ECC51009.2020.9143962> (2020).
- Yang, X., Ding, F., Zhang, D. & Zhang, M. Vehicular trajectory big data: Driving behavior recognition algorithm based on deep learning. In *Artificial Intelligence and Security, Communications in Computer and Information Science* (eds Sun, X. *et al.*) 324–336 (Springer, 2020). https://doi.org/10.1007/978-981-15-8086-4_30.
- Wang, Y., Wei, L. & Chen, P. Trajectory reconstruction for freeway traffic mixed with human-driven vehicles and connected and automated vehicles. *Transp. Res. Part C Emerg. Technol.* **111**, 135–155. <https://doi.org/10.1016/j.trc.2019.12.002> (2020).
- Barmounakis, E. & Geroliminis, N. On the new era of urban traffic monitoring with massive drone data: The pNEUMA large-scale field experiment. *Transp. Res. Part C Emerg. Technol.* **111**, 50–71. <https://doi.org/10.1016/j.trc.2019.11.023> (2020).
- Krajewski, R., Bock, J., Kloeker, L. & Eckstein, L. The highD Dataset: A drone dataset of naturalistic vehicle trajectories on German highways for validation of highly automated driving systems. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)* 2118–2125. <https://doi.org/10.1109/ITSC.2018.8569552> (2018). ISSN: 2153-0017.
- Makridis, M., Mattas, K., Anesiadou, A. & Ciuffo, B. OpenACC. An open database of car-following experiments to study the properties of commercial ACC systems. *Transp. Res. Part C Emerg. Technol.* **125**, 103047. <https://doi.org/10.1016/j.trc.2021.103047> (2021).
- Knoop, V. L. *et al.* Platoon of SAE Level-2 automated vehicles on public roads: Setup, traffic interactions, and stability. *Transp. Res. Rec.* **2673**, 311–322. <https://doi.org/10.1177/0361198119845885> (2019).
- Tian, J. *et al.* On the role of speed adaptation and spacing indifference in traffic instability: Evidence from car-following experiments and its stochastic model. *Transp. Res. Part B Methodol.* **129**, 334–350. <https://doi.org/10.1016/j.trb.2019.09.014> (2019).
- Toledo, T., Koutsopoulos, H. N. & Ahmed, K. I. Estimation of vehicle trajectories with locally weighted regression. *Transp. Res. Rec.* **1999**, 161–169. <https://doi.org/10.3141/1999-17> (2007).
- Montanino, M. & Punzo, V. Trajectory data reconstruction and simulation-based validation against macroscopic traffic patterns. *Transp. Res. Part B Methodol.* **80**, 82–106. <https://doi.org/10.1016/j.trb.2015.06.010> (2015).
- Fard, M. R., Shariat Mohaymany, A. & Shahri, M. A new methodology for vehicle trajectory reconstruction based on wavelet analysis. *Transp. Res. Part C Emerg. Technol.* **74**, 150–167. <https://doi.org/10.1016/j.trc.2016.11.010> (2017).
- Punzo, V., Formisano, D. J. & Torrieri, V. Nonstationary Kalman filter for estimation of accurate and consistent car-following data. *Transp. Res. Rec.* **1934**, 2–12. <https://doi.org/10.1177/0361198105193400101> (2005).
- Chen, X., Li, Z., Yang, Y., Qi, L. & Ke, R. High-resolution vehicle trajectory extraction and denoising from aerial videos. *IEEE Trans. Intell. Transp. Syst.* **22**, 3190–3202. <https://doi.org/10.1109/ITTS.2020.3003782> (2021).
- Treiber, M. & Kesting, A. Microscopic calibration and validation of car-following models—A systematic approach. *Procedia Soc. Behav. Sci.* **80**, 922–939. <https://doi.org/10.1016/j.sbspro.2013.05.050> (2013).
- Makridis, M., Fontaras, G., Ciuffo, B. & Mattas, K. MFC free-flow model: Introducing vehicle dynamics in microsimulation. *Transp. Res. Rec.* **2673**, 762–777. <https://doi.org/10.1177/0361198119838515> (2019).
- Fadhloun, K. & Rakha, H. A novel vehicle dynamics and human behavior car-following model: Model development and preliminary testing. *Int. J. Transp. Sci. Technol.* **9**, 14–28. <https://doi.org/10.1016/j.ijtst.2019.05.004> (2020).
- Axenie, C. *et al.* Fuzzy modeling and inference for physics-aware road vehicle driver behavior model calibration. <https://doi.org/10.2139/ssrn.4195952> (2022).
- Raissi, M., Perdikaris, P. & Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* **378**, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045> (2019).

29. Karniadakis, G. E. *et al.* Physics-informed machine learning. *Nat. Rev. Phys.* **3**, 422–440. <https://doi.org/10.1038/s42254-021-00314-5> (2021).
30. Shi, R., Mo, Z., Huang, K., Di, X. & Du, Q. A physics-informed deep learning paradigm for traffic state and fundamental diagram estimation. *IEEE Trans. Intell. Transp. Syst.* **23**, 11688–11698. <https://doi.org/10.1109/TITS.2021.3106259> (2022).
31. Thodi, B. T., Khan, Z. S., Jabari, S. E. & Menéndez, M. Incorporating kinematic wave theory into a deep learning method for high-resolution traffic speed estimation. *IEEE Trans. Intell. Transp. Syst.* **23**, 17849–17862. <https://doi.org/10.1109/TITS.2022.3157439> (2022).
32. Mo, Z., Shi, R. & Di, X. A physics-informed deep learning paradigm for car-following models. *Transp. Res. Part C Emerg. Technol.* **130**, 103240. <https://doi.org/10.1016/j.trc.2021.103240> (2021).
33. Ciuffo, B., Makridis, M., Toledo, T. & Fontaras, G. Capability of current car-following models to reproduce vehicle free-flow acceleration dynamics. *IEEE Trans. Intell. Transp. Syst.* **19**, 3594–3603. <https://doi.org/10.1109/TITS.2018.2866271> (2018).
34. Makridis, M. A., Anesiadou, A., Mattas, K., Fontaras, G. & Ciuffo, B. Characterising driver heterogeneity within stochastic traffic simulation. *Transportmetrica B Transp. Dyn.* <https://doi.org/10.1080/21680566.2022.2125458> (2022).
35. Commission, E. Vehicle categories (2016).
36. He, Y. *et al.* Introducing electrified vehicle dynamics in traffic simulation. *Transp. Res. Rec.* **2674**, 776–791. <https://doi.org/10.1177/0361198120931842> (2020).
37. Chao, Q. *et al.* A calibrated force-based model for mixed traffic simulation. *IEEE Trans. Vis. Comput. Graph.* <https://doi.org/10.1109/TVCG.2021.3128286> (2021).
38. Thiel, C., Schmidt, J., Van Zyl, A. & Schmid, E. Cost and well-to-wheel implications of the vehicle fleet CO₂ emission regulation in the European Union. *Transp. Res. Part A Policy Pract.* **63**, 25–42. <https://doi.org/10.1016/j.tra.2014.02.018> (2014).
39. Staacks, S., Hütz, S., Heinke, H. & Stampfer, C. Advanced tools for smartphone-based experiments: Phyphox. *Phys. Educ.* **53**, 045009. <https://doi.org/10.1088/1361-6552/aac05e> (2018).

Acknowledgements

We want to thank Mr. Panagiotis Kalaitzidis for his help with data gathering concerning the second campaign and Prof. Francesco Corman for the insightful discussion and comments on our work. This work was partly supported by the Swiss National Science Foundation (SNSF) under the project RECCE, “Real-time traffic estimation and control in a connected environment”, contract No. 200021-188622.

Author contributions

M.M. conceptualization, data analysis, visualization, methodological approach, draft, and final manuscript version. A.K. methodological approach, draft, and final manuscript version. M.M. is the corresponding author.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-023-28202-1>.

Correspondence and requests for materials should be addressed to M.A.M.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher’s note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023