



HAL
open science

Towards a robust and consistent estimation of a vehicle's mass

Mathieu Randon, Benjamin Quost, Nassim Boudaoud, Dirk von Wissel

► **To cite this version:**

Mathieu Randon, Benjamin Quost, Nassim Boudaoud, Dirk von Wissel. Towards a robust and consistent estimation of a vehicle's mass. The European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML/PKDD 2020), Sep 2020, Ghent, Belgium. hal-03122765

HAL Id: hal-03122765

<https://hal.science/hal-03122765v1>

Submitted on 1 Feb 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Towards a robust and consistent estimation of a vehicle’s mass

Mathieu Randon^{1,2,3,4}, Benjamin Quost^{1,3}, Nassim Boudaoud^{2,3}, and Dirk von Wissel⁴

¹ Heudiasyc laboratory, UMR UTC-CNRS 7253

² Roberval laboratory, FRE UTC-CNRS 2012

³ Sorbonne Universités, Université de Technologie de Compiègne

⁴ Renault S.A.S

Abstract. A detailed knowledge of a vehicle’s characteristics makes it possible to monitor its dynamic behavior, energy consumption, and wear. In this preliminary study, we address the problem of learning a robust and consistent mass estimator from data provided by embedded sensors, which are subject to uncertainties. Consistency refers to the ability to comply with physical laws — Newton’s second law in the case of mass estimation, robustness to the capacity to infer from uncertain or scarce data. This preliminary work aims at defining the problem and providing some guidelines with respect to constructing a robust and consistent mass estimator from uncertain data. Simple experiments on a Renault vehicle confirm the feasibility and the interest of learning a consistent vehicle model so as to increase the estimation accuracy of vehicle consumption.

Keywords: robustness and consistency · physics-constrained estimation · dynamical mechanical system modeling · informed machine learning

1 Introduction

In the automotive industry, vehicle modeling has become a key issue for various purposes, such as consumption estimation and optimization. A vehicle model heavily depends on a number of parameters which determine the vehicle’s behavior. For instance, the load, which may vary according to its usage, influences its dynamics (and therefore its energy consumption) as well as its wear. On-board measurement of the vehicle mass is thus a key issue towards a better monitoring and control of vehicle performances.

Measuring the vehicle mass classically requires the use of sensors, the accuracy of which directly depends on their price. An alternative approach consists in constructing estimators (also referred to as observers, inferential sensors, or virtual sensors) [15,10], which spares making these additional expenses, but may result in poor parameter estimates. Some physical values describing the vehicle’s behavior cannot be directly measured, and thus require the use of such estimators. Obviously, the quality and quantity of available data then inevitably restricts the accuracy of the resulting estimations.

This preliminary study principally aims at clearly setting the problem, and in particular justify the need for robust and consistent estimators. Robustness directly refers to the ability to provide accurate estimates in presence of uncertain or scarce input data. Our purpose is to develop a vehicle model using machine learning, through which data uncertainty can be propagated. Thus, the resulting parameter estimates will depend on data quality. Consistency, on the other hand, is defined as the ability of the model to comply with the well-established laws of physics — here, Newton’s laws of motion — to which the system is subject. The use of physical knowledge to constrain the training of a vehicle model is related to the field known as informed machine learning [21], and theory-guided data science [9]. Our ultimate goal is to be able to exploit consistency in order to reduce the uncertainty arising from the data, while remaining robust to noisy or scarce data.

A short list of references focusing on mass estimation is given in Section 2. Section 3 discusses the problem of building a robust consistent vehicle mass estimator via machine learning. In Section 4, we report some preliminary experiments realized on data measured on a real vehicle. Finally, Section 5 concludes the paper and provides several future research directions.

2 Existing works

Several articles have investigated estimating a vehicle’s mass from data. In 2005, [26] considered a longitudinal dynamic vehicle model, and proposed an online mass estimation strategy using recursive least square. Road slope was also jointly estimated as it is highly coupled to mass in the dynamic model considered. Other contributions considered using extended Kalman filters [24], sensor-based road slope estimation [6], Lyapunov-based observers [13,12], and road slope transition filtering [11].

More recently, several studies addressed the problem of robustness. For instance, [1] proposed to increase robustness in Kalman and recursive least squares filtering approaches, by adding a conditioning term and managing outliers. Alternatively, [25] proposed to estimate mass and road slope directly using a neural network. In [23], uncertainty is managed using a particle filtering approach, whereas [17] uses maximum likelihood estimation jointly with polynomial chaos theory in order to recursively estimate a reduced-order state-space vehicle model.

3 Problem statement

Sensors are subject to uncertainties, which may come from various sources (aleatoric perturbations due to their functioning, biases due to inappropriate locations in the vehicle, limited scope of use). As a consequence, estimators are themselves subject to the uncertainties coming from their input data (which may be scarce or noisy), to the parameters on which they depend, and due to the assumptions underlying the model of the estimator itself (which reflect a particular context of usage).

We aim at providing estimators which are robust to low-quality data, by exploiting the physical principles to which the system modeled is subject — in our case, the vehicle should satisfy Newton’s laws of motion. The requirement that the estimator be consistent, i.e. that it satisfies these physical principles, will help guiding the estimation process towards a robust solution.

3.1 Robust inference

One of the difficulties to proceed with mass estimation comes from the high uncertainty of the data gathered via the controller area network (CAN) embedded into a series production vehicle. A now classical distinction can be made between aleatoric and epistemic uncertainty [22,7].

In our problem, aleatoric uncertainty will systematically pervade the data at hand — for example, aleatoric delays generated by the CAN protocol (typically, $50 \pm 50\text{ms}$), or noise in the sensors’ measurement process, due to their physical conception (for instance, a same acceleration will deform the silica spring of the accelerometer with an aleatoric variation due to small frictions, elasticity, or geometry variations in sensor physics). Classically, this uncertainty is quantified via physical measurement approaches.

Epistemic uncertainty, on the other hand, reflects the lack of knowledge with respect to the actual vehicle’s behavior in terms of quantity and quality of data. First of all, despite the increasing number of embedded sensors in modern vehicles, the three dimensional dynamic of this multi-body system remains a very complex model [27], which would require a large quantity of data provided by high quality sensors to be properly assessed. For instance, a sensor on the vehicle damper would provide a quite accurate estimation of the vehicle mass by measuring the damper travel; however, this sensor being expensive, it cannot be made available in series vehicles. As a consequence, it is in practice very difficult to gather a large number of high-quality data about the vehicle’s behavior in each of its operating modes. Another source of epistemic uncertainty corresponds to the biases, unique to each vehicle, caused by the manufacturing or assembly processes. For instance, an internal Renault study showed that the braking torque data were subject to such uncertainties, since the embedded black box sensor and does not take into account the brake setting nor the wear.

Note that a third source of uncertainty may be identified with the choice of the vehicle model, which may be too simple or inappropriate. For instance, the use of a nominal mass value for estimating the parameters in a hybrid model, such as mentioned in Section 3.3, falls into this category, since subsequent measurements or predictions realized during the vehicle’s life (e.g. related to consumption) may be inaccurate. Model assumptions may also be restrictive: for instance, damping is often neglected, which prevents to accurately model vehicle pitching; friction forces of the wheel on the road are considered independent of velocity, which is notoriously wrong (friction forces are equal to zero when the vehicle is motionless on a flat road); the acceleration is assumed to be measured at the moving center of gravity, which may not always be the case in practice.

Being able to correctly model the vehicle’s dynamic while being robust to poor quality data motivates the use of advanced machine learning methods of a hybrid model, where physics principles make it possible to guide the estimation of the statistical model. In addition, it seems necessary to propagate the uncertainty through the model, so as to assess to which extent it is appropriate given the data at hand and the physics principles involved. The interest of such a “self-aware” [22] estimator of the vehicle mass is twofold. First of all, should the mass estimate be used in subsequent estimates (for instance of the vehicle consumption), its validity could be taken into account. Besides, it would ideally provide an indicator of the vehicle’s operating modes in which either the data and the model are in conflict, which would indicate poor inappropriate model assumptions; or for which the estimates are too uncertain, thus pointing out the necessity to gather additional data in order to better describe the vehicle behavior.

3.2 Consistency

Due to the data being scarce or poor in some operating modes of the vehicle, statistical approaches to modeling the vehicle behavior are prone to overfitting. Our proposal to address this undesirable outcome consists in constraining the model to satisfy some constraints, imposed by the physics laws to which the vehicle is subject. In the present case, as any dynamical system, the moving vehicle should satisfy Newton’s second law of motion [14]:

$$\frac{d}{dt} (m(t) \cdot \vec{v}(t)) = \sum \vec{f}_{ext}(t); \quad (1)$$

here, $m(t)$ is the total mass of the vehicle frame and its load, $\vec{v}(t)$ is the speed of the center of gravity of the vehicle frame in the Galilean referential, and $\sum \vec{f}_{ext}(t)$ refers to the sum of all external forces on the vehicle frame at time t .

We will make the following additional assumptions: the vehicle mass m is constant over a trip; the vehicle frame is a rigid body, and its center of gravity is constant; and the external forces are composed of both forces generated by the gravity g which are proportional to m , and forces which are independent of m :

$$\sum \vec{f}_{ext}(t) = m \cdot \sum \vec{f}_g(t) + \sum \vec{f}_g(t).$$

Then, Equation (1) can be projected onto a direction \vec{i} , which yields:

$$m \cdot a(t) = f(t) + \varepsilon_{gen}(m, t) \quad (2)$$

with

$$a(t) = \left(\frac{d}{dt} \vec{v}(t) + \sum \vec{f}_g(t) \right) \cdot \vec{i}, \quad f(t) = \sum \vec{f}_g(t) \cdot \vec{i}.$$

The term a relates to the sum of vehicle frame kinematic and gravity effects on the vehicle frame, whereas f relates to the sum of mechanical actions applied on

the vehicle frame which are independent of gravity; finally, $\varepsilon_{gen}(m, t)$ represents possible errors due to the assumptions underlying the generic model which are considered to depend on m . For example, the assumption that the mass m is constant over a trip amounts to neglect weight loss due to fuel consumption (this can leave to an error of about 30kg) or weight variations due to the arrival or the leave of passengers.

3.3 Consistent online mass estimation

A straightforward approach to mass estimation would lead to invert equation 2. Obviously, due to measurement uncertainties, this approach requires solving an online least-squares optimization problem:

$$\hat{m} = \arg \min_m \sum_{t \in \mathcal{T}} (m \cdot a(t) - f(t))^2, \quad (3)$$

where \mathcal{T} is the set of timestamps of the portion of the trip used for the estimation, during which mass variations are neglected. The solution \hat{m} to Equation (3) is the slope of the “best-fit” line in the (a, f) referential (see Figure 2). To proceed with this estimation, it is necessary to obtain values for $a(t)$ and $f(t)$ from the CAN embedded into the vehicle.

Physical model The physical model, which is the most frequently found in the literature, defines $a(t)$ and $f(t)$ using the mechanical knowledge of the system along the longitudinal axis \vec{x} of the vehicle frame:

$$\begin{cases} f_{phys}(t; \theta_f) = \theta_0 + \theta_1 \cdot t_w(t) + \theta_2 \cdot \ddot{x}(t) + \theta_3 \cdot \dot{x}(t)^2 + \varepsilon_f(t) \\ a_{phys}(t; \theta_a) = \ddot{x}(t) + g \sin(\alpha(t)) + \theta_4 \cdot g \cos(\alpha(t)) + \varepsilon_a(t) \end{cases}, \quad (4)$$

where $\dot{x}(t)$ is the vehicle speed, $\ddot{x}(t)$ is the vehicle acceleration (derivative of $\dot{x}(t)$), $t_w(t)$ is the traction torque on the vehicle wheels, $\alpha(t)$ is the slope of the road, g is the gravity acceleration, θ are the model parameters, and $\varepsilon_f(t)$ and $\varepsilon_a(t)$ represents errors made because of the model assumptions.

Statistical model It might be tempting to use a purely statistical model of the vehicle so as to estimate a and f . However, this approach did not prove to be relevant so far.

Indeed, even high-quality sensors (which cannot be embedded on production models, for obvious economical reasons) are subject to uncertainties, due to drift, suboptimal location in the vehicle, or restricted scope of validity. The collection of a large dataset of a high quality is therefore very expensive and time-consuming. Experiments showed that using machine learning without any additional background knowledge performs poorly, mainly due to overfitting, as will be seen in Section 4.

Hybrid model A hybrid model combines a statistical approach with a physical component, which can either be a physical model such as described above, or a set of constraints arising from the dynamics of the vehicle. The interest is to avoid making prior assumptions about a physical model which are oversimplistic or have a restricted range of validity. It allows for a good compromise between model adaptability (provided by machine learning) and universality (guaranteed by satisfying the laws of physics). As an example, if $f(t)$ (the sum of mechanical forces independent of gravity) is to be statistically estimated from data whereas $a(t)$ is determined as previously, the physical model (4) results in the following hybrid model:

$$\begin{cases} f_{stat}(t; \theta_f) = h_f(y_f(t), \theta_f) + \varepsilon_f(t) \\ a_{phys}(t; \theta_a) = \ddot{x}(t) + g \sin(\alpha(t)) + \theta_A \cdot g \cos(\alpha(t)) + \varepsilon_A(t) \end{cases}, \quad (5)$$

where $y_f(t)$ are the input data of the statistical model h_f obtained by the CAN, $\alpha(t)$ is the slope of the road, and θ_f is the set of parameters governing this statistical model.

The learning task consists in computing the best estimate $\hat{\theta}$ of $\theta = (\theta_a, \theta_f)$ using a set of learning samples:

$$\hat{\theta} = \arg \min_{\theta} j(\theta) = \sum_{t \in \mathcal{T}} (m_0 \cdot a_{\theta_a}(t) - f_{\theta_f}(t))^2, \quad (6)$$

under the constraints imposed by Equation 2 and both models assumptions:

$$a_{\theta_a}(t) = h_a(y_a(t), \theta_a) + \varepsilon_a(t), \quad f_{\theta_f}(t) = h_f(y_f(t), \theta_f) + \varepsilon_f(t).$$

Note that computing $\hat{\theta}$ requires to know the mass of the vehicle: hence, it is performed during the vehicle's conception step using a nominal mass m_0 . As a consequence, the estimated components $a(\cdot; \hat{\theta}_a)$ and $f(\cdot; \hat{\theta}_f)$ may only be valid, during the vehicle's use, if its actual mass is close to m_0 .

Multiple methods have been recently developed to proceed with such a physics-guided learning strategy [19,3,28].

4 Preliminary experiments on a real case

The problem presented in Section 3 has been applied to a 1300 kg passenger vehicle. The global dataset contains 63h54min of CAN data recorded at 10Hz, on open road, for 5 different additional load cases (90kg, 155kg, 255kg, 325kg and 375kg). This dataset is split into three parts: 67% training data, 8% for validation, and 25% in the test set. The training data corresponds to the first samples of the global trip (roughly, from $t = 1$ to $t = 17000$ in Figures 3 and 5), and to vehicle load weights of 155kg, 255kg, and 325kg (as well as a few instances with an additional load of 375kg). The test data (from $t = 17000$ on) correspond to additional loads of 375kg and 90kg.

4.1 Data preprocessing

Raw data are filtered with a 3Hz low pass filter (Shannon criteria) and then 0.3Hz to remove high frequency noises. The road slope (α) is calculated using a sensor-based method [18,6].

The physical model (4) is valid only under several conditions: straight position, stabilized slope, no braking, no clutching, moderate speed, and moderate longitudinal accelerations. Consequently, all data samples which do not satisfy those conditions may actually be considered as noisy, since they would induce a bias in the estimation process. These data can therefore be removed from the dataset — which leaves us with 1% of the data.

Feature engineering, which corresponds to transforms of the input data based on mechanical knowledge, is performed to reduce the number of layers needed for the neural network. For instance, VS^2 (square of vehicle speed) and V_{1000} are calculated from other input data. The former allows to model aerodynamic resistance at high speed, the latter represents the gearbox ratio which impacts friction forces of the transmission. Estimates of derivatives with respect to time are commonly used to provide information about the temporal dynamics of the vehicle. These transformations may however increase the level of uncertainty in the training data. Note that the features which appear in the statistical model h_f have been selected using the mechanical knowledge of the system only.

4.2 Training

The approach described in Section 3.3 is applied on the data. The model h_f is a feed-forward neural network displayed in Figure 1. Its accuracy is assessed via curves displaying the mass estimate during time, the distribution of residuals (Figure 4), as well as the determination coefficient

$$R^2 = 1 - \frac{\sum_{\forall t} (m(t) - \hat{m}(t))^2}{\sum_{\forall t} (m(t) - \bar{m})^2}.$$

Note that this coefficient can take negative values when computed over a set of data distinct from the training sample.

Both the purely physical model and the hybrid model allow for different levels of load to be distinguished: for each cloud of sample data corresponding to a particular load, a specific slope can be estimated (see Figure 2). After validation, the hybrid model obtains an overall R^2 score of 0.95, which is slightly higher than that obtained with the physical model (0.92).

4.3 Results on mass estimation

The online estimation of vehicle mass (3.3) is performed recursive least squares [8] with a forgetting factor of 0.957 (10 minutes driving in average depending on contexts). Results are shown in Figures 3, 4, and 5.

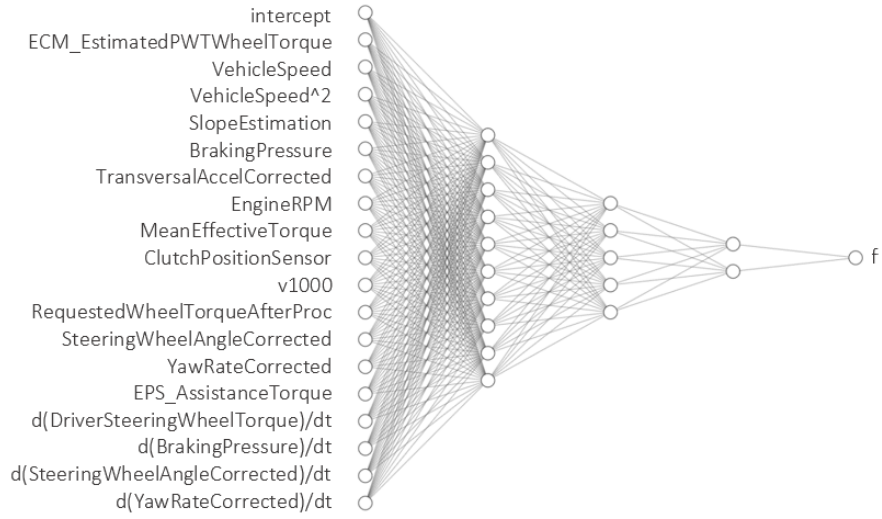


Fig. 1. Chosen feed-forward neural network architecture: 19 inputs, 3 layers of 10, 5 and 2 neurons with rectified linear unit activation function.

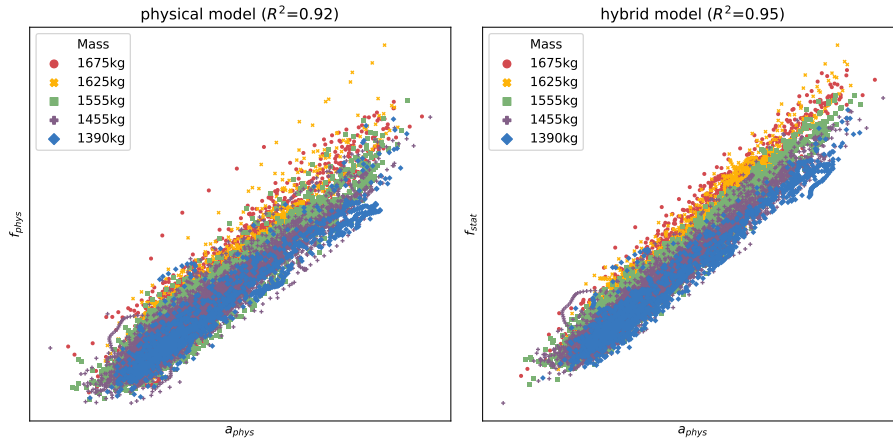


Fig. 2. Vehicle data projected in the (a, f) regression referential learnt via the physical model (left, $R^2=0.92$) and the hybrid model (right, $R^2=0.95$).

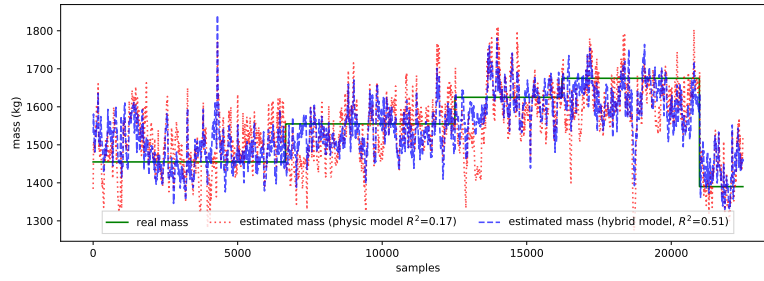


Fig. 3. Mass estimated with the physical model (red dotted line, $R^2=0.17$) and the hybrid model (blue line, $R^2=0.51$), both after a selection of the training data.

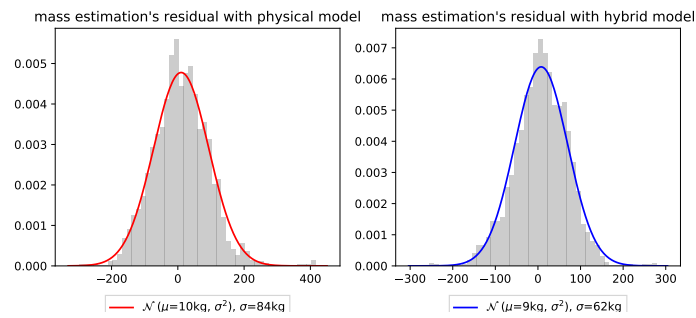


Fig. 4. Distribution of residuals for the physical model (red line, $\sigma=84\text{kg}$) and the hybrid model (blue line, $\sigma=62\text{kg}$) both after a selection of the training data.

We first comment the results obtained after a selection of the training data based on the driving conditions (Figure 3). Overall, the mass is estimated with a 95% confidence interval accuracy of $[-119kg, +137kg]$ (see Figure 4). The hybrid model significantly improves accuracy (R^2 score of 0.51) compared to the purely physical model (R^2 score of 0.17). However, the errors may be large in some driving contexts, and a more stable estimation can be obtained with an increase of the forgetting factor.

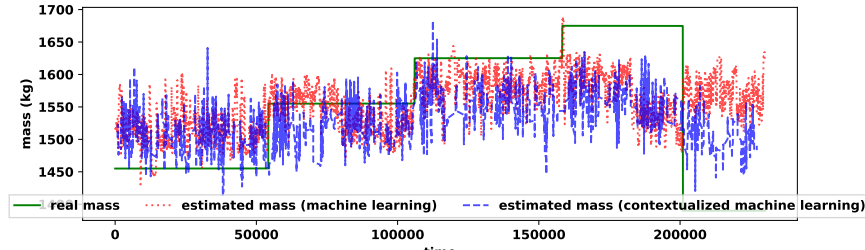


Fig. 5. Mass estimated with a purely statistical model, without training data selection (red dotted line, $R^2=0.04$) and after training data selection (blue line, $R^2=0.21$).

Should an unconstrained statistical model be directly estimated — i.e., the mass is estimated using a neural network without any physical constraints, the resulting model now appears to be unable to relate mass to the input data, as can be seen in Figure 5. The estimation seems to be only marginally affected by load variations of the vehicle. Selecting the training instances based on driving conditions appeared to improve the results only in a limited way (overall R^2 score of 0.04 without selection, against 0.21 with selection). This illustrates the need of data selection and the sensitivity of pure statistical approaches to overfitting.

Note that although the R^2 score might seem better for the unconstrained statistical model with selected examples than for the physical model, both of the unconstrained approaches actually perform poorly (R^2 scores of -0.28 without sample selection, and 0.06 with sample selection) in the test sequence (from $t = 17000$ on), compared to the physical and hybrid models. Besides, we found out that selecting the training instances may significantly improve the results, but may also lead to overfit some specific driving contexts if the screening conditions are too restrictive, most probably due to the training set being then depleted. The use of a statistical approaches is more adaptable to a specific vehicle configuration and its embedded sensors. However, collecting an exhaustive dataset with loading cases in diverse driving conditions is limited by industrial resources, therefore pure machine learning approach trends to overfit. We expect a careful uncertainty quantification and propagation to help detecting those driving contexts for which too few data are available.

Last, if the estimation accuracy can appear to be poor compared to other results found in literature in Section 2, this is essentially due to poor quality and quantity of available sensors to describe the complexity of a real vehicle dynamics. However, even if the 95% confidence interval accuracy is quite limited (standard deviation $\sigma \simeq 62\text{kg}$), it remains sufficient to distinguish an overload use of the vehicle from normal load contexts. This distinction can be already be used onboard to improve control strategy of powertrain, steering, braking and damping systems of the vehicle. This is all the more true for commercial vehicles, for which the load can reach more than 100% of the empty vehicle’s weight.

5 Conclusions and perspectives

In this preliminary contribution, we presented a problem where some parameters influencing a dynamical system are to be estimated. More particularly, we focused on a vehicle where the mass is to be determined from measurements provided by embedded sensors. Important choices influencing data quality and quantity are guided by industrial requirements. These constraints, together with the high influence of the driving context on the vehicle’s behavior, strongly motivates the use of cautious models, where the uncertainty coming from the quality of the data at hand, systematic errors made by the sensors, and oversimplified physical models are propagated to the model output. This seems even more necessary as the global model output is to be used in a more general vehicle model. The objective is to find methods to build a reliable estimator from uncertain or scarce data that can be generalized to the estimation of other first-order parameters of a vehicle.

For this purpose, two families of approaches can be deployed. In a purely physical model, the dynamical system is described using the mechanical knowledge of the system only. Alternatively, a purely statistical approach usually more adaptable to infer the dynamic behavior directly from sensors available in the vehicle. In the case of mass estimation, both purely statistical, and purely physical approaches are limited, the former by epistemic uncertainties (representing lack of quality and quantity of data), and the latter by its range of validity. We therefore advocate a hybrid approach, where some of the physical components are replaced by a statistical counterpart. This makes it possible to avoid restricting the model to over-simplistic physical assumptions, while reducing trends of machine learning to overfit with lack of data, by incorporating additional knowledge (laws of dynamics to which the vehicle is subject) to the training process.

Preliminary results obtained in Section 4 demonstrate the interest of the hybrid approach: constraining the statistical component makes it possible to increase the mass estimation accuracy. This shows that using a complex statistical model makes it possible to account for complex dynamic behaviors which are difficult to explicitly assess using parameterized dynamics equations. The results obtained motivate requiring that the statistical model be consistent with respect to the laws of dynamics. Feature selection and model selection (based on statistical tests) may allow for further improvements. It could also be interesting to

apply other models found in the literature 2 on the proposed dataset to obtain an objective comparison.

Future work will be conducted in several directions. First of all, we only assessed the interest of taking into account data uncertainty in the crudest of ways — i.e., by selecting data according to their quality and relevance to the vehicle’s operating mode. This makes it possible to lower the bias induced by atypical data, at the price of a loss of information. We will therefore focus on uncertainty propagation. More particularly, we will investigate generalized least squares, Bayesian or robust Bayesian approaches [10], as well as imprecise-probabilistic and credal formalizations, in order to propagate data scarcity or low quality to the estimates, and to take into account the degree of conflict between the input data and the vehicle model estimated. For Bayesian approaches, weak prior knowledge on the vehicle load could be derived from the number of passengers (given by the number of fastened seat belts) and the remaining fuel in the vehicle. However, this would be restricted to passenger cars, and would exclude commercial vehicles. The results may be compared to those previously obtained using generative models [5], maximum likelihood estimation with polynomial chaos modeling [16], or Gaussian process regression via Kriging [2,20].

In addition, this approach may be generalized to estimating several first-order parameters such as aerodynamic resistance, battery state of health, tire efficiency, or driving conditions [4]. In this more general setting where several parameters are to be jointly estimated, an additional challenge comes from the interdependencies between these parameters. Hopefully, consistency may then help proceeding with the estimation in a large number of operating contexts while preventing overfitting.

Acknowledgments

The authors are grateful to Renault teams for providing the dataset used in this paper and for sharing feedback on internal mass estimation algorithms and their limits.

References

1. Altmannshofer, S., Endisch, C.: Robust vehicle mass and driving resistance estimation. vol. 2016 American Control Conference (ACC), pp. 6869–6874 (2016)
2. Castric, S., Denis-Vidal, L., Cherfi, Z., Blanchard, G.J., Boudaoud, N.: Modeling pollutant emissions of Diesel engine based on Kriging models: a comparison between geostatistic and Gaussian Process approach. IFAC Proceedings Volumes **45**, 1708–1715 (2012)
3. Chen, R.T.Q., Rubanova, Y., Bettencourt, J., Duvenaud, D.: Neural Ordinary Differential Equations. Advances in Neural Information Processing Systems **31**, 6571–6583 (2017)
4. Fontaras, G., Zacharof, N.G., Ciuffo, B.: Fuel consumption and CO2 emissions from passenger cars in Europe – laboratory versus real-world emissions. Progress in Energy and Combustion Science **60**, 97 – 131 (2017)

5. Grathwohl, W., Chen, R.T.Q., Bettencourt, J., Sutskever, I., Duvenaud, D.K.: FFJORD: Free-form Continuous Dynamics for Scalable Reversible Generative Models. ArXiv [abs/1810.01367](#), 1–13 (2019)
6. H., I.K., Kim, Bang, J., Huh, K.: Development of estimation algorithms for vehicle's mass and road grade. *International Journal of Automotive Technology* **14**, 889–895 (2013)
7. Hüllermeier, E., Waegeman, W.: Aleatoric and epistemic uncertainty in machine learning: An introduction to concepts and methods. ArXiv [abs/1910.09457](#), 1–17 (2019)
8. Ioannou, P., Fidan, B.: *Adaptive Control Tutorial*. Society for Industrial and Applied Mathematics (2006)
9. Karpatne, A., Atluri, G., Faghmous, J.H., Steinbach, M.S., Banerjee, A., Ganguly, A.R., Shekhar, S., Samatova, N.F., Kumar, V.: Theory-Guided Data Science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering* **29**, 2318–2331 (2017)
10. Khatibisepehr, S., Huang, B., Khare, S.: Design of inferential sensors in the process industry: A review of Bayesian methods. *Journal of Process Control* **23**, 1575–1596 (2013)
11. Kidambi, N., Harne, R., Fujii, Y., Pietron, G., Wang, K.: Methods in vehicle mass and road grade estimation. *SAE International Journal of Passenger Cars - Mechanical Systems* **7**, 11 (2014)
12. Mahyuddin, M., Na, J., Herrmann, G., Ren, X., Barber, P.: Adaptive observer-based parameter estimation with application to road gradient and vehicle mass estimation. *IEEE Transactions on Industrial Electronics* **61**, 2851–2863 (2014)
13. McIntyre, M., Ghotikar, T., Vahidi, A., Song, X., Dawson, D.: A two-stage Lyapunov-based estimator for estimation of vehicle mass and road grade. *IEEE Transactions on Vehicular Technology* **58**, 3177–3185 (2009)
14. Newton, I.: *Philosophiæ Naturalis Principia Mathematica* (1686)
15. Patwardhan, S.C., Narasimhan, S., Jagadeesan, P., Gopaluni, B., Shah, S.L.: Non-linear bayesian state estimation: A review of recent developments. *Control Engineering Practice* **20**, 933–953 (2012)
16. Pence, B., Fathy, H., Stein, J.: Recursive maximum likelihood parameter estimation for state space systems using polynomial chaos theory. *Automatica* **47**, 2420–2424 (2011)
17. Pence, B., Fathy, H., Stein, J.: Recursive estimation for reduced-order state-space models using polynomial chaos theory applied to vehicle mass estimation. *IEEE Transactions on Control Systems Technology* **22**, 224–229 (2014)
18. Planche, G., Verbeke, B., Zaoui, K.: Device and method for estimating the charge of a motor vehicle (2013)
19. 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. *Journal of Computational Physics* **378**, 686–707 (2019)
20. Rasmussen, C.E.: *Gaussian Processes for Machine Learning*, pp. 63–71. The MIT Press (2006)
21. von Rüden, L., Mayer, S., Beckh, K., Georgiev, B., Giesselbach, S., Heese, R., Kirsch, B., Pfrommer, J., Pick, A., Ramamurthy, R., Walczak, M., Garcke, J., Bauckhage, C., Schücker, J.: Informed machine learning – a taxonomy and survey of integrating knowledge into learning systems. ArXiv [abs/1903.12394](#), 1–20 (2019)

22. Senge, R., Bösner, S., Dembczyński, K., Haasenritter, J., Hirsch, O., Donner-Banzhoff, N., Hüllermeier, E.: Reliable classification: Learning classifiers that distinguish aleatoric and epistemic uncertainty. *Information Sciences* **255**, 16–29 (2014)
23. Sun, S., Zhang, N., Walker, P., Lin, C.: Intelligent estimation for electric vehicle mass with unknown uncertainties based on particle filter. *IET Intelligent Transport Systems* **14**, 463–467 (2020)
24. Sun, Y., Li, L., Yan, B., Yang, C., Tang, G.: A hybrid algorithm combining EKF and RLS in synchronous estimation of road grade and vehicle mass for a hybrid electric bus. *Mechanical Systems and Signal Processing* **68-69**, 416–430 (2016)
25. Torabi, S., Wahde, M., Hartono, P.: Road grade and vehicle mass estimation for heavy-duty vehicles using feedforward neural networks. vol. 1, pp. 316–321 (2019)
26. Vahidi, A., Stefanopoulou, A., Peng, H.: Recursive least squares with forgetting for online estimation of vehicle mass and road grade: Theory and experiments. *Vehicle System Dynamics* **43**, 31–55 (2005)
27. Venture, G., Ripert, P.J., Khalil, W., Gautier, M., Bodson, P.: Modeling and identification of passenger car dynamics using robotics formalism. *IEEE Transactions on Intelligent Transportation Systems* **7**, 349–359 (2006)
28. Zhu, Y., Zabaras, N., Koutsourelakis, P.S., Perdikaris, P.: Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data. *Journal of Computational Physics* **394**, 56–81 (2019)