Data-driven on-line load monitoring in garbage trucks

Valentina Breschi *, **Simone Formentin** *, **Davide Todeschini** *,
Alberto L. Cologni **, **Sergio M. Savaresi** *

* *Dipartimento di Elettronica e Informazione, Politecnico di Milano,
Piazza Leonardo da Vinci 32, 20133 - Milan, Italy.*
** *E-Novia SpA, Via San Martino 12, 20122 - Milan, Italy.*

Abstract: The payload of garbage trucks may vary substantially over the time, affecting both the vehicle performance and driving safety. Information on the load in real-time could thus play a key role for monitoring and diagnostics. Unfortunately, physical sensors directly measuring the vehicle mass are usually too costly for commercial trucks, while the correlation between consecutive values of the load is not considered by most of existing approaches for mass estimation. Since this correlation characterizes load variations in garbage trucks, this paper proposes an ad-hoc approach for payload estimation, which relies on inertial sensors only. To minimize the tuning effort, we introduce a strategy to automatically select the key tunable parameters of the estimator. The effectiveness of the proposed approach is demonstrated on experimental data collected on a real truck.

1. INTRODUCTION

Several automotive applications, including safety-oriented controls and fault detection systems, require an accurate knowledge of the vehicle mass to perform properly. Since the mass of a heavy-duty vehicle can vary as much as 400% depending on its payload (Kidambi et al. (2014)), retrieving reliable information on the load in real-time is of paramount importance to optimize the performance of the vehicle and detect potential safety hazards. A correct estimation of the payload can also help improving fleet management strategies (Stefansson and Lumsden (2008)), ultimately leading to potential savings.

Multiple approaches have been already proposed to tackle the problem of mass estimation (Kidambi et al. (2014)), which are mainly model-based. These methods can be divided into strategies relying on the model for the lateral dynamics of the vehicle, *e.g.,* Wenzel et al. (2004), and the ones based on the model for its longitudinal dynamics, as Fathy et al. (2008); Vahidi et al. (2003a,b); Winstead and Kolmanovsky (2005). Available approaches for load estimation can be further distinguished as methods that return an estimate of the load regularly over time, and techniques providing load estimates only when specific events are detected. Even though most of these approaches require a dedicated sensor setup, which can be fairly expensive, some methods do exist (see Reineh et al. (2014)) that require a set of (low-cost) inertial sensors only. However, due to their dependence on a mathematical model for the vehicle dynamics, they might perform poorly when the modelling assumptions are not satisfied. Additionally,

Fig. 1. Experimental setup. The red circles indicate the positions of the inertial sensors.

existing approaches for load estimation usually rely on the hypothesis of constant/slowly varying loads, which is quite unrealistic when considering garbage trucks. These issues are bypassed by the frequency-based approach in Bottelli et al. (2014) that relies on the *offline* identification of a *black-box* model describing load variations, obtained as a combination of spectral features of the available signals that can be linked to variations in the payload. This approach has proven to enable the user to accurately detect mass variations, while relying on a simple sensor setup and operations that can be easily performed on-board of commercial vehicles. Indeed, this approach requires the use of a single *Inertial Measurement Unit* (IMU), which can be integrated into garbage trucks with ease.

In this paper, we focus on real-time payload estimation of garbage trucks, with the ultimate goal of reducing their operating costs while detecting potentially unsafe working conditions. Indeed, real-time estimation of the load allows the driver to be notified before the vehicle is overloaded, thus avoiding possible structural damages and circumventing fees related to excessive payloads, see Italian Highway Code (1992). At the same time, by monitoring load variations on-line, one can actively detect crossings of
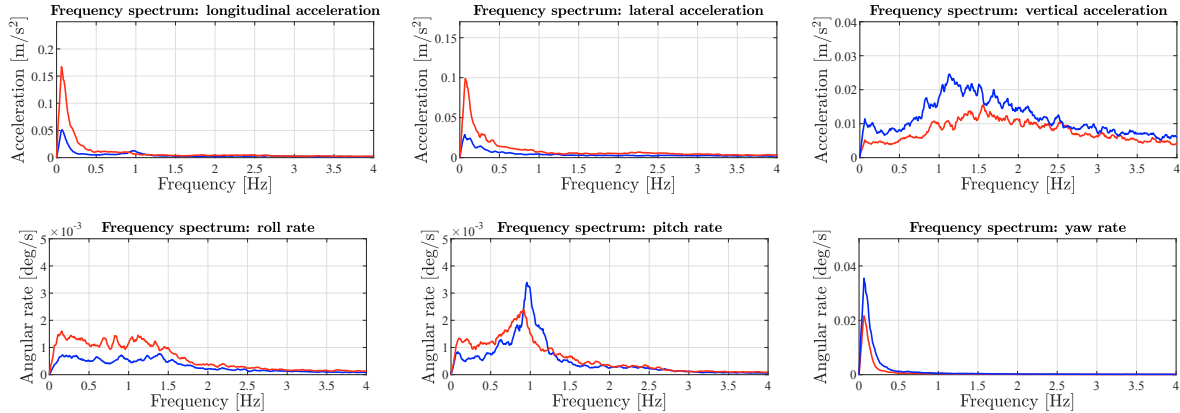
Fig. 2. Examples of low-frequency spectra of accelerations and angular rates. First (blue) and last (red) 15 minutes of acquisition.

critical thresholds, thus retrieving useful information for diagnostic/prognostic purposes. Since the payload might considerably vary each time the garbage truck stops to collect waste, the approach in Bottelli et al. (2014) seems the most suited for the problem at hand. Therefore, we propose a method based on the detection of changes in the frequency spectra to track variations in the payload but, differently from Bottelli et al. (2014), we further exploit the characteristics of the application at hand by imposing a correlation between consecutive load estimates. Moreover, as we rely on the sensor setup depicted in Fig. 1, our vehicle is equipped with two IMUs. Therefore, unlike in Bottelli et al. (2014), we can exploit more information, at the price of additional efforts required to select relevant signals and fuse the available information.

The paper is organized as follows. In Section 2, we describe the experimental setup and discuss the results of a preliminary data analysis. The approach for load estimation is presented in Section 3, while an empirical solution to reduce the tuning effort is discussed in Section 4. Experimental results are presented in Section 5, and the paper is ended by some concluding remarks.

2. EXPERIMENTAL SETUP AND DATA ANALYSIS

As shown in Fig. 1, we consider a garbage truck equipped with two IMUs mounted on the rear axle and the dump body of the truck, respectively. Both IMUs consist of three accelerometers and three gyroscopes, all with sampling rate at 100 Hz. The truck is further equipped with a *Global Positioning System* (GPS), that returns the position of the vehicle every 0.02 s. Since it is unlikely that even at the same location the amount of collected waste is always the same, *e.g.,* because of seasonality, it is reasonable to assume that GPS measurements provide little/no insights on the current payload. Therefore, GPS measurements are discarded, allowing us to neglect the difference in the sampling rates. Therefore, the only data that can be used to estimate the load are the accelerations and the angular rates returned by the IMUs. As in Bottelli et al. (2014), we analyze these signals by considering their frequency spectra during the first and the last 15 minutes of acquisition, so as to understand if they are affected by variations in the payload. Since each acquisition terminates once the truck is unloaded, these subsets of data are related to the *beginning* and the *end* of the time interval in which the
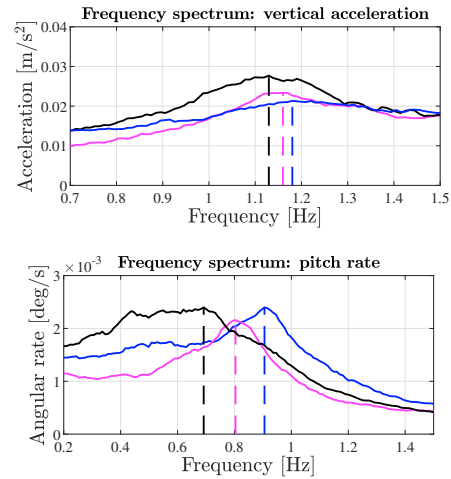

Fig. 3. Examples of vertical acceleration and pitch rate low-frequency spectra obtained by splitting one dataset into 5 equally long subsets. First (blue), third (magenta) and last (black) segments. The dashed vertical lines indicate the positions of the peaks.

garbage truck is in use and, thus, they can be associated with the conditions of *almost empty* and *full* truck. Since we are evaluating the differences in the spectra for two opposite load configurations, our analysis is expected to highlight which signals are actually influenced by variations in the load, so to eventually reduce the amount of data to be fused. As shown in Fig. 2, changes in the spectra are particularly evident at low frequencies, where the magnitude of the spectral peaks clearly changes as a function of the load. Moreover, the frequency of the spectral peaks shifts based on the payload for the vertical acceleration and the pitch rate only. Provided this additional property of these signals, it seems reasonable to estimate the load by exploiting vertical accelerations and pitch rates only. Note that a similar conclusion is also reached by the authors in Bottelli et al. (2014), hence endorsing our choice. As we aim at retrieving an estimate of the load in real-time, the vertical acceleration and the pitch rate are further analyzed to detect changes in the spectral peaks throughout an acquisition. To this end, we split each dataset into 5 consecutive subsets of equal length, which are associated with increasing values of the load. Then, we compute the spectrum for each subset and we evaluate the magnitude

Fig. 4. Interpolating linear function (black) vs ratio index/load pairs (blue). The ratio indexes are computed for the pitch rate measured by the IMU mounted on the rear axle of the truck.

and position of the spectral peaks with respect to prior subsets. As a result, we observe a behavior similar to the one shown in Fig. 3. This highlights that the position of the spectral peaks does change according to variations in the load.

3. DATA-DRIVEN LOAD ESTIMATION

Since the results reported in Section 2 are consistent with the ones in Bottelli et al. (2014), the estimator is designed by following a procedure similar to the one considered there. Specifically, the vertical accelerations and pitch rates are preprocessed to solely retain the information that is useful for load estimation. After this preprocessing phase, a *black-box* model is then learned, which returns the value of the load for a given set of selected features.

Since load variations mainly affect the spectral peaks at low frequencies, we extrapolate the information at these frequencies by constructing the *ratio index*, a dimensionless indicator already considered in Bottelli et al. (2014) that provides information on the shape of the spectra. For a given signal y, this index is computed as the following power spectral ratio

$$R_y = \frac{P_{y_1}}{P_{y_2}}, \tag{1}$$

with P_{y_1} and P_{y_2} being the spectral powers of the signals y_1 and y_2, that are obtained by filtering y with two distinct *bandpass filters* $F_1(s, f_{l_1}, f_{u_1})$ and $F_2(s, f_{l_2}, f_{u_2})$ [1]. The cutoff frequencies $\{f_{l_i}\}_{i=1}^2$ and $\{f_{u_i}\}_{i=1}^2$ are tunable parameters, with $f_{l_i} < f_{u_i}$ for $i = 1, 2$. It is worth pointing out that a proper tuning of these cutoff frequencies is of paramount importance to accurately reconstruct the load. Since our ultimate goal is to estimate the load in real-time, the ratio indexes have to be computed iteratively. In the training phase, they can be obtained from data acquired between consecutive measurements of the load. However, when the estimator is deployed the load is not measured and, thus, the ratio indexes can only be computed through data collected over consecutive intervals of predefined length.

As we have no prior knowledge on the relationship between the load and the indexes, we consider a general estimator belonging to the class:

$$\mathcal{M}_{nl} : \quad \ell(n, \theta) = g(\theta), \tag{2}$$

where $n \in \mathbb{N}$ is a counter whose value increases every time the load is estimated. Both the parameter vector $\theta \in \mathbb{R}^{n_\theta}$ and $g : \mathbb{R}^{n_\theta} \to \mathbb{R}$ in (2) have to be retrieved from data,

[1] s denotes the Laplace variable.

since they are *unknown*. In particular, $g(\theta)$ is generally a nonlinear function of the unknown parameters θ and the information available up to the instant when the n-th estimate of the load is computed, *e.g.,* current and past ratio indexes and either past measurements or estimates of the load.

To select the structure of the estimator, we have preliminarily analyzed the relationship between the ratio indexes and the measurements of the load on the available dataset. The results of this investigation are similar to the ones shown in Fig. 4, leading us to the conclusion that the relation between the ratio indexes and the load is approximately linear. As a consequence, we identify an affine model by modifying (2) as follows

$$\mathcal{M} : \quad \ell(n, \theta) = \tilde{X}(n)'\theta, \tag{3}$$

with $\tilde{X}(n) \in \mathbb{R}^{n_\theta}$ defined as

$$\tilde{X}(n) = \begin{bmatrix} X(n) \\ 1 \end{bmatrix}, \tag{4}$$

and $X(n)$ being a properly defined feature vector. Given a training set $\{X(n), L(n)\}_{n=1}^N$ comprising N features/load pairs, the parameter vector $\theta \in \mathbb{R}^{n_\theta}$ is identified by minimizing the quadratic cost function

$$\mathcal{J}_\theta = \sum_{n=1}^N (L(n) - \ell(n, \theta))^2, \tag{5}$$

with respect to θ, through standard least-squares.

Inspired by Bottelli et al. (2014), we initially learn a *static* affine model by using the feature vector $X(n)$:

$$X_s(n) = \begin{bmatrix} R_{a_z,1}(n) R_{a_z,2}(n) R_{g_y,1}(n) R_{g_y,2}(n) \end{bmatrix}'. \tag{6}$$

Once the parameter vector θ is learned, the estimate of the load is obtained via a linear combination of the ratio indexes of the vertical accelerometers and the pitch gyros, respectively denoted as $R_{a_z,i}$ and $R_{g_y,i}$, with $i = 1, 2$. This choice entails that no priority is attributed to the data collected from either of the two available sensors.

Although the *static* model (3)-(6) exploits the results of our analysis in Section 2, it does not leverage a distinctive feature of the considered problem, namely the *correlation* between consecutive values of the load. This additional information is accounted for by introducing an extended linear model, with feature vector defined as

$$X(n) = \begin{bmatrix} \ell(n-1) \ \ell(n, \hat{\theta}_s)' \end{bmatrix}', \tag{7}$$

with $\ell(n, \hat{\theta}_s)$ being the estimate obtained by using the static model. The new estimator resulting from (3)+(7) accounts for the correlation between successive values of the load, while fusing the information provided by the available sensors as the *static* estimator. Note that, given the load estimated via the static model, the unknown parameters of the *correlation-based* estimator (3)+(7) are still identified by minimizing the cost function in (5) with respect to a new set of parameters $\theta \in \mathbb{R}^2$. This *dynamic* model can then be used for real-time load estimation by exploiting previous estimates of the payload and the current estimate of the payload obtained with the static estimator.

4. AN AUTOMATIC TUNING PROCEDURE

The computation of the ratio indexes requires the design of two different bandpass filters for each signal of interest,

Fig. 5. Distribution of the peaks' frequencies for the first and last 90000 samples acquired during each round of the garbage truck.


Fig. 6. Filters width vs ratio index/load pairs (blue) and their interpolating function (black). The ratio indexes are computed for the pitch rate measured by the IMU mounted on the rear axle of the truck.

so to isolate the frequency ranges in which the spectra are more *sensitive* to load variations. This entails the selection of 4 cutoff frequencies for each pair of filters. Given the results of the analysis in Section 2, in our setting we need to design *eight* different filters, namely to tune 16 parameters to isolate frequencies around the spectral peaks.

For the ratio indexes to be informative, the frequency ranges passed by $F_1(s, f_{l_1}, f_{u_1})$ and $F_2(s, f_{l_2}, f_{u_2})$ should not overlap, so that the indexes combine different spectral information. To account for this requirement, while reducing the number of parameters to be tuned, we impose the following constraint:

$$f_{u_1} = f_{l_2}. \qquad (8)$$

As the upper cutoff frequency of F_1 is equal to the lower cutoff frequency of F_2, this guarantees that the ranges of each pair of filters do not overlap.

Although it is crucial for the filters not to overlap, it is equally important that the ratio indexes embed all information that are useful for load estimation. Based on

Table 1. Auto-tuned parameters of the band-pass filters.

	$f_{l_1,(\cdot)}$	$f_{u_1,(\cdot)} = f_{l_2,(\cdot)}$	$f_{u_2,(\cdot)}$
Vertical accelerometers	1.00 Hz	1.22 Hz	1.44 Hz
Pitch Gyros	0.72 Hz	0.81 Hz	1.00 Hz

the results of Section 2, this implies that the filters must isolate the frequencies around spectral peaks. A possible approach to fulfill this requirement, is proposed in Bottelli et al. (2014) and it can be easily modified to fit our setting. However, this greedy method requires the evaluation of a cost function for different choices of the unknown cutoff frequencies until satisfactory results are attained. This implies a *time consuming tuning phase*, to exhaustively explore all the possible combinations of parameters. At the same time, this procedure might lead to *unreliable choices* for the cutoff frequencies, if the space of all the possible bandpass filters is not thoroughly searched. To overcome these limitations, we introduce an alternative tuning procedure that leverages on the results reported in Section 2 and, in particular, on the fact that variations in the payload lead to changes of the spectra in proximity of the spectral peaks. Since this phenomena was also observed by the authors in Bottelli et al. (2014), we expect the proposed method to be fairly general.

As an initial step, we select the frequencies already constrained in (8) as follows:

$$f_{u_1} = f_{l_2} = \mu, \qquad (9)$$

where μ is the sample mean of the frequencies associated with the spectral peaks. The remaining parameters are then tuned based on the dispersion of the peaks' frequencies above and below the mean value μ. Specifically, f_{l_1} and f_{u_2} are chosen as

$$f_{l_1} = \mu + D_\mu(f), \qquad (10a)$$
$$f_{u_2} = \mu - d_\mu(f), \qquad (10b)$$

where $D_\mu(f)$ and $d_\mu(f)$ are the maximum absolute distances of the peaks' frequencies from μ above ($f > \mu$) and below ($f < \mu$) the mean value itself, respectively. Note that it is generally advisable to compute $D_\mu(f)$ and $d_\mu(f)$ neglecting peaks that are too close to the mean value, while trimming the tails of the peaks distribution.

By considering the specific experimental setting considered in the paper, the analysis of the available dataset has resulted in the distributions of the peaks' frequencies reported in Fig. 5. It is clear that the peaks always lie within a fairly limited range of frequencies, which can be easily isolated by selecting the bandpass filters as proposed above. We have also noticed that the use of filters of increasing width results in a loss of linearity in the relation between the ratio indexes and the payload, as shown in Fig. 6. The proposed tuning strategy should allow us to trade off between two conflicting objectives, namely reducing the width of the filters as much as possible for the linear model to be representative of the actual features/load relationship, while spanning a sufficiently broad range of frequencies, not to filter out information that might be useful.

5. EXPERIMENTAL RESULTS

The available experimental dataset consists of data collected throughout 19 days, whose features are detailed

Table 2. Features of the validation sets.

Dataset	Mean load μ_L	Maximum Load L_{max}
#1	854.7 kg	1915.0 kg
#2	1875.5 kg	4000.0 kg
#3	3651.7 kg	7560.0 kg

Table 3. Reduced (5 parameters) vs complete (13 parameters) static models.

|	RMSE [kg]	
Dataset	Complete	Reduced
#1	518.1	422.3
#2	683.4	655.4
#3	524.2	511.2

Fig. 7. Distribution of time intervals between consecutive load acquisitions.

in Table 2. Both static and correlation-based models are trained by using sixteen out of 19 dataset, so that the training set is sufficiently heterogeneous. The remaining 3 datasets are used to validate the estimators. These sets are chosen so that they are sufficiently diverse, therefore allowing us to assess the performance of the estimators in different settings of practical interest. As shown in Table 2, we test the estimator for "light" (dataset #1), "intermediate" (dataset #2) and "heavy" (dataset #3) loads. In the training phase, the ratio indexes are computed with data collected between two consecutive measurements of the load. On the other hand, in practice the estimator is not aware of the instants when the payload is eventually evaluated. Therefore, in validation the load is estimated with a preselected rate. By looking at the distribution in Fig. 7, the load is reconstructed every 10 minutes. Since in validation any information on the load is used only as ground-truth to assess the quality of our estimates, when using the correlation-based estimator we construct the regressor by replacing $\ell(n-1)$ with the estimated load. For the estimates to be coherent with the physical meaning of the unknowns, the estimated load is set to zero every time it is negative. By analyzing our datasets further, we have also noticed that the average frequency of the spectral peaks is around 1.20 Hz for both vertical accelerometers, while it is around 0.85 Hz for the pitch gyros. Therefore, it seems reasonable to design two distinct pairs of filters only, customized for the vertical accelerometers and the pitch gyros respectively, to compute the spectral indexes required to construct the regressors in equations (6) and (7). This leads to a significant reduction in the number of tunable parameters, that is now equal to *eight*. The remaining cutoff frequencies are selected via the tuning procedure presented in Section 4, whose results are reported in Table 1.

The performance of the estimators are quantitatively assessed by using the Root Mean Square Error (RMSE) [kg], that is defined as


Fig. 8. Measured payload (blue) vs load estimated every 10 minutes with the static (black) and the correlation-based (red) estimators.

Table 4. Static vs correlation-based model.

	RMSE [kg]	
Dataset	Static	Correlation-based
#1	422.3	230.5
#2	655.4	495.5
#3	511.2	357.6

Table 5. Estimator in Bottelli et al. (2014) vs correlation-based estimator.

	RMSE [kg]	
Dataset	Bottelli et al. (2014)	Correlation-based
#1	422.8	230.5
#2	702.5	495.5
#3	515.4	357.6

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{n=1}^{N}(L(n) - \ell(n,\hat{\theta}))^2}, \qquad (11)$$

with accurate estimates corresponding to small values for the RMSE. Notice that to compute this quality index, the estimated load has to be compared with the measured one. For this comparison to be fair, the load cannot be estimated every 10 minutes, but it has to be computed at the same unevenly spaced instants at which the actual payload is measured. Accordingly, only when evaluating these indexes we exploit this additional information on the load.

To validate the results of the preliminary analysis in Section 2, *i.e.,* to show that the only signals that are significant for load estimation are the vertical accelerations and the pitch rates, we compare the performance of two static models identified. One model is trained by exploiting all 12 available signals, while the other one is obtained by using the vertical accelerations and the pitch rates only. As shown in Table 3, the use of more signals leads to a slight deterioration on the performance of the *static* estimator. This phenomenon can be related with the increased complexity of the model to be trained, combined with the use of data that are less relevant to the considered estimation task. The static estimator is then compared with the correlation-based one. As shown in Table 4 and in Fig. 8, the introduction of the correlation between consecutive values of the load enhances the quality of the estimator and it leads to smoother estimates. It is worth pointing out that the RMSEs are computed with respect to estimates obtained at the same unevenly spaced instants at which the actual load is measured, while the results reported in Fig. 8 are obtained by reconstructing the load every 10 minutes. The performance of the correlation-based estimator is also compared with the one of the

(a) Validation set: #1

(b) Validation set: #2
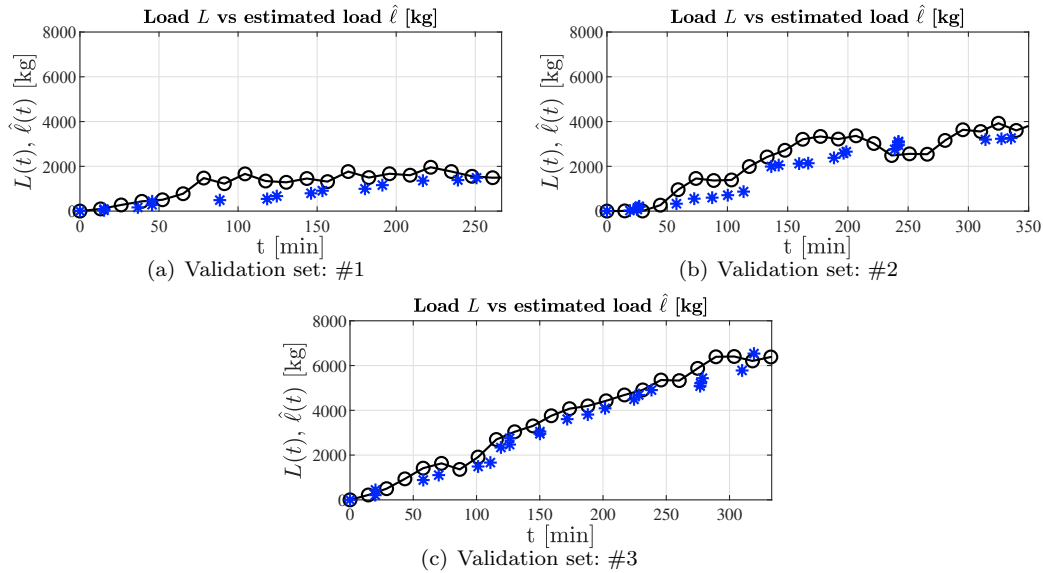
(c) Validation set: #3

Fig. 9. Measured payload (blue) vs load estimated every 10 minutes with the correlation-based estimator (black).

approach proposed in Bottelli et al. (2014). To adhere with the framework presented in Bottelli et al. (2014), the load is estimated exploiting only the vertical acceleration and the pitch rate collected from the IMU mounted on the rear axle. Differently from Bottelli et al. (2014), we estimate the load by using an affine model, as the nonlinear one proposed in Bottelli et al. (2014) performs poorly in our setting. The RMSEs obtained with the two estimators are reported in Table 5, proving that the introduction of the correlation between consecutive estimates and the use of the information provided by an additional IMU enhance the quality of the estimate. In Fig.9 we finally show the estimates obtained with the correlation-based estimator for all validation datasets, along with the actual values of the payload. Also in this case the estimate in retrieved every 10 minutes by simulating the correlation-based model in open loop, while the measurements of the load are available at unevenly spaced instants. These results show that the payload reconstructed with the correlation-based estimator follows the actual variations in the load, with an average error of approximately 12%. Moreover, Fig.9 indicates that the estimator performs better for the third dataset, suggesting the potential need for multiple estimators to handle different load configurations.

6. CONCLUSIONS

We have presented a frequency-based method to estimate the load of a garbage truck. Inspired by the seminal work in Bottelli et al. (2014), we resort to the estimation of a model describing load variations, obtained by fusing the information collected from two IMUs mounted on the truck. Nonetheless, the proposed approach exploits on the main feature of the considered problem, namely the correlation between consecutive values of the load.
Future research includes a refinement of the model to explicitly account for the monotonic increase of load.

REFERENCES

Bottelli, S., M., T., Boniolo, I., and Savaresi, S. (2014). Online estimation of vehicle load and mass distribution for ground vehicles. *IFAC Proceedings Volumes*, 47(3), 8439 – 8444.

Fathy, H., Kang, D., and Stein, J. (2008). Online vehicle mass estimation using recursive least squares and supervisory data extraction. In *2008 American Control Conference*, 1842–1848.

Italian Highway Code (1992). *art. 162, Legislative Decree 285/1992*.

Kidambi, N., Harne, R., Fujii, Y., Pietron, G., and Wang, K. (2014). Methods in vehicle mass and road grade estimation. *SAE Int. J. Passeng. Cars - Mech. Syst.*, 7, 981–991.

Reineh, M., Enqvist, M., and Gustafsson, F. (2014). IMU-based vehicle load estimation under normal driving conditions. In *53rd IEEE Conference on Decision and Control*, 3376–3381.

Stefansson, G. and Lumsden, K. (2008). Performance issues of smart transportation management systems. *International Journal of Productivity and Performance Management*, 58(1), 55–70.

Vahidi, A., Druzhinina, M., Stefanopoulou, A., and Peng, H. (2003a). Simultaneous mass and time-varying grade estimation for heavy-duty vehicles. In *Proceedings of the 2003 American Control Conference, 2003.*, volume 6, 4951–4956.

Vahidi, A., Stefanopoulou, A., and Peng, H. (2003b). Experiments for online estimation of heavy vehicles mass and time-varying road grade. In *Proceedings of the 2003 ASME International Mechanical Engineering Congress and Exposition.*, 451–458.

Wenzel, T., Burnham, K., Blundell, M., and Williams, R. (2004). Approach to vehicle state and parameter estimation using extended Kalman filtering. In *Proceedings of the International Symposium on Advanced Vehicle Control (AVEC)*.

Winstead, V. and Kolmanovsky, I. (2005). Estimation of road grade and vehicle mass via model predictive control. In *Proceedings of 2005 IEEE Conference on Control Applications, 2005. CCA 2005.*, 1588–1593.