Strain, T., Gunner, S., & Wilson, E. (2019). Estimation of Vehicle Counts from the Structural Response of a Bridge. In *2019 International Conference on Smart Infrastructure and Construction (ICSIC 2019)* Thomas Telford (ICE Publishing). https://doi.org/10.1680/icsic.64669.751

## University of Bristol - Explore Bristol Research
### General rights

International Conference on Smart Infrastructure and Construction 2019 (ICSIC):
Driving data-informed decision-making

ice Publishing

CSIC Cambridge Centre for
Smart Infrastructure
& Construction

# ESTIMATION OF VEHICLE COUNTS FROM THE STRUCTURAL RESPONSE OF A BRIDGE

*T.J. Strain[*1], S. Gunner[2] and R.E. Wilson[3]*

[1,2,3]*Department of Engineering Mathematics, University of Bristol, Bristol, United Kingdom*
*\* Corresponding author*

**ABSTRACT** In this paper an offline method for counting vehicles travelling across a bridge through its structural response under loading is developed. Readings from a single vertically oriented accelerometer fixed to the bridge are normalised and then summarised by instantaneous amplitude envelopes. The envelopes for a single vehicle have a profile that is log-normal in appearance. Least squares fitting is used in conjunction with the Akaike information criterion to fit the envelope time series as the superposition of $n$ log-normal functions (other functional forms are also considered). It is hypothesised that this fit describes $n$ vehicles travelling across the bridge. This method is applied to data from a previous study on rapid deployment of structural health monitoring systems undertaken on the Clifton Suspension Bridge (CSB) in Bristol. The data from a single strain gauge-based accelerometer installed as part of a wireless sensor network (WSN) on the bridge is used and demonstrates the value added by this method to pre-existing asset monitoring systems. A prediction accuracy of 74% is achieved on a labelled test set.

## 1. Introduction

Structural health monitoring (SHM) systems enable those responsible for the safe and proper maintenance of critical infrastructure to make better data-driven decisions in order to ensure that damage and deterioration is properly managed. Sensors are fitted to an asset so that the displacement and acceleration of its structural components can be recorded. A great deal of work in applying novel sensing capabilities is undertaken in this area. Chan et al. (2006) installed Fiber Bragg grating sensors to take strain measurements of the Tsing Ma bridge in Hong Kong. Their optics-based sensing capability provided comparable strain measurements to a previous study. Oh et al. (2015) implemented a vision-based system to monitor the acceleration of a building from a multi-camera installation. Their system captured identical behaviour to traditional accelerometers. To monitor the entirety of an asset, sensors are often deployed at various locations as part of a wireless sensor network (WSN), such as a WSN of video cameras to monitor the structural health of a bridge (Basharat, et al., 2005).

Sensing capabilities like these provide rich data sources to perform remote and intelligent SHM. However, what is less well researched is using the information from these sensors for a secondary purpose. Once a sensing capability has been installed onto an asset for SHM, there are no other costs besides the required maintenance of the WSN, and there is great potential to add value by using the data for other purposes.

This paper describes a method that reuses data from accelerometers installed as part of a WSN on the Clifton Suspension Bridge (CSB) in Bristol. Originally installed as a feasibility study in deploying SHM systems on the bridge, the sensors monitored the acceleration and displacement of the CSB over a two-month period. In this paper, a method is developed to count the number of vehicles that cross the bridge from the accelerometer data. Traditionally inductance loops are used to count vehicles, however, they are expensive to install and require the road to be closed and dug up. A method that can count vehicles by using already available data might therefore represent savings both in cost and in time.

## 2. Background

In January 2017 a WSN was fitted to the CSB for a study in rapid deployment of SHM systems (Gunner, et al., 2017). From 25 January to 16 March 2017, the bridge's structural response was captured and stored so that any changes in the bridge's physical condition might be monitored.

The CSB stands at 75m above the River Avon and spans 214m across the Avon Gorge joining Clifton to Leigh Woods. The deck is suspended via 81 pairs of vertical iron rods. Since its opening in 1864, the bridge has remained an important piece of infrastructure, carrying three million vehicles every year (Yeung & Smith, 2005). The bridge has one lane in each direction that is accessible to cyclists. A 20mph speed limit is in force on the bridge and so it takes a minimum of 24 seconds to drive across the entire span. There is a footpath on each side of the bridge for pedestrians.

Figure 1 Sensor deployment in the CSB trial. The Clifton end is in the foreground and the Leigh Woods end is furthest away from the viewer. The camera is pointing in a broadly westerly direction.
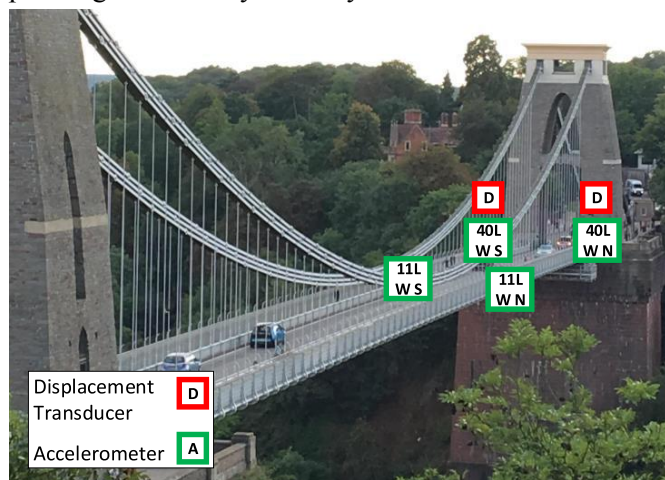


A toll barrier at each end of the CSB records when a vehicle enters the bridge, truncated to the nearest minute. Vehicles travelling towards Leigh Woods pass through the toll barrier at the Clifton end of the bridge and vice versa. There is no information available on when cyclists and pedestrians enter or leave the bridge.

Since a large number of vehicles use the CSB (over 8,000 per day on average), the toll barrier data, together with the structural responses captured in the study, provide a rich data set to help develop a vehicle count estimation method. Due to its age and design, the CSB is more compliant under loading than modern bridges. Hence the deflections due to vehicles tend to be larger and more easily identifiable.

*2.1 Sensor deployment*

Due to time and cost constraints, just four accelerometers were deployed on the bridge, in two pairs. Each pair straddles the width of the bridge deck, with one accelerometer on the north side and one on the south side, so that both vertical and torsional displacement can be measured. Each accelerometer is identified by the vertical iron suspension rod it is closest to, counting from the middle of the span, and whether it is on the north or south (N or S) side of the bridge, see Figure 1. The accelerometers have a range of 10ms$^{-2}$, a sampling rate of 4,000 readings per minute and a resolution of 5.2×10$^{-5}$ms$^{-2}$. The sampling rate which is relatively low, is chosen to prolong battery life, whilst resolving the principle modes of the bridge.

The first accelerometer pair (11LW N and 11LW S) is placed 26.8m from the midpoint, since Macdonald (2008) showed that this is the optimal location to pick up all of the deck's natural frequencies up to 3Hz.

The second pair (40 LW N and 40 LW S) is placed towards the Leigh Woods end of the bridge, where the structural response to a vehicle entering or exiting there is greatest.

Finally, in addition to the accelerometers, displacement transducers were deployed at the north and south sides of the deck at the Leigh Woods end. The transducers measure vertical displacement, however, their data is not used in this study. A comprehensive description of the sensor deployment is provided in Gunner et al. (2017).

*2.2 Accelerometer windows*

Figure 2 shows two-minute windows of the bridge's structural response to vehicles travelling across the bridge in each direction, where the directions are determined by the toll barrier data. The figure shows that accelerometers 11LW N and 11LW S are featureless under loading for both directions of travel, whereas a clear response is observed in accelerometers 40LW N and 40LW S. There is no obvious difference between the responses recorded by the 40LW N and 40LW S accelerometers, therefore we will only consider the acceleration recorded by 40LW N in this paper, which we denote as $a_t$ at time $t$.

# 3. Data processing

*3.1 Segmentation and matching of accelerometer data*

In order to develop a method to count vehicles, the accelerometer data is segmented into windows with corresponding ground truth vehicle counts provided by the toll barrier data.

The toll barriers record when a vehicle enters the bridge truncated to the nearest minute. That is, a vehicle arriving at time $t \coloneqq$ YYYYDDMM hh:mm:ss is attributed to

$$T(t) \coloneqq \text{YYYYDDMM hh:mm:00.} \qquad (1)$$

Minute by minute, we thus have counts for the number of vehicles that enter each end of the bridge. The accelerometer data $a_t$ can then be segmented into the corresponding one-minute windows that begin at time $T(t)$.

Consider the following situation. A vehicle enters the bridge at time $t_1$, and is still travelling across the bridge at time $t_2$ such that $T(t_2) > T(t_1)$. In this case, the vehicle may cause a structural response in the one-minute window after that in which it enters the bridge.

To ensure that all structural responses of a vehicle entering in a given minute are captured, we group that minute's toll barrier data with the accelerometer data for that and the following minute, resulting in a two-minute window.

Further, to prevent confusion with the structural response caused by vehicles that enter the bridge either before or after the selected minute, we take forward for analysis only those

minutes for which the immediately preceding and succeeding toll barrier counts are zero. The accelerometer has a sampling interval of $\Delta t = 0.015s$, and so each two-minute window of accelerometer data can be expressed as $a_i$, for $i = 1,2,\ldots,m$ where $m = 8,000$.

Figure 2 Panels show accelerometer readings for two-minute windows where respectively (a) two vehicles enter from the Leigh Woods end and (b) a single vehicle enters from the Clifton end. The structural responses to each vehicle are especially clear in these examples, however, they are not as clear in other windows.



(a)



(b)
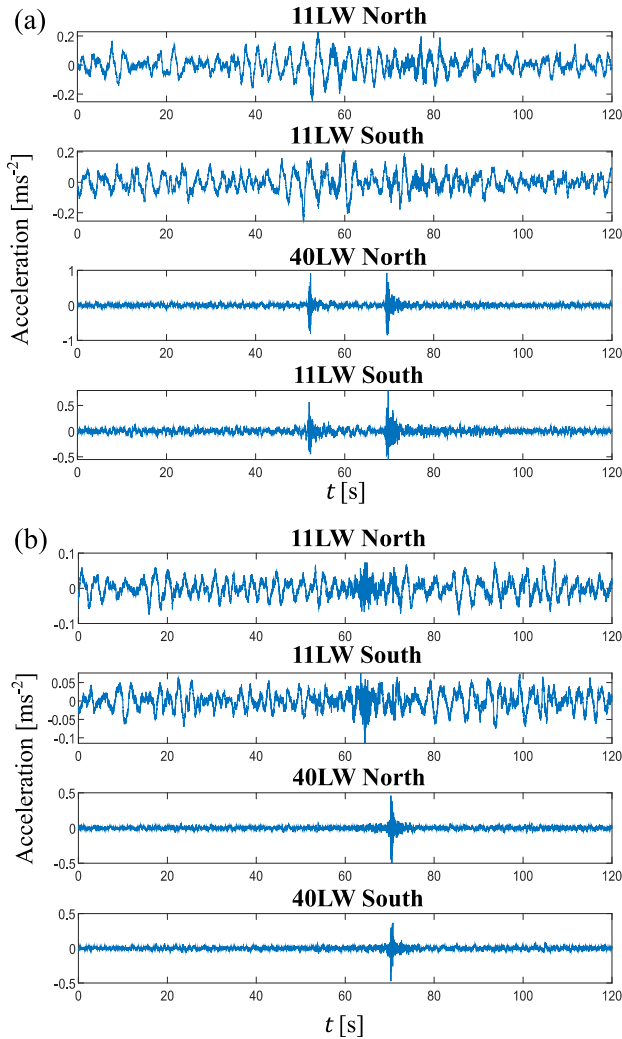
### 3.2 Instantaneous amplitude envelope computation

To ensure that each window of accelerometer data has a similar minimum and maximum amplitude, a normalisation is performed on each window in the form

$$\tilde{a}_i = 2\left[\frac{a_i - \min(a_i)}{\max(a_i) - \min(a_i)}\right] - 1, \qquad (2)$$

to give normalised readings $\tilde{a}_i$ between -1 and 1.

We now summarise accelerometer data in a smoothed form known as the instantaneous amplitude envelope (IAE) (Picinbono, 1997). We consider the three following methods for computing the IAE, see Figure 3.

1. **Peak to Peak.** The accelerometer data is replaced by a cubic spline that interpolates local maxima, provided they are separated by at least $d$ time-steps (Biswas & Si, 2011). Here $d = 5$ was observed to achieve the required level of smoothing, but the computed envelope is not necessarily positive, so the method is discarded.

2. **Analytic Signal.** The IAE is found from the magnitude of the discrete-time analytic signal (Reilly, et al., 1994). The computed envelopes are less smooth than we require, so the method is discarded.

3. **Moving Root Mean Square (MRMS).** MRMS computes the root mean square from a window of width $2w$ that slides across the accelerometer data (Albu & Heydt, 2003). Thus smoothed readings $\hat{a}_i$ are found by

$$\hat{a}_i = \sqrt{\frac{1}{2w_i - 1}\sum_{i-w_i}^{i+w_i} \tilde{a}_i}, \qquad (3)$$

where $w_i = \min(i, w, m - i)$. The computed envelope is thus strictly positive, and a choice of $w = 20$ (window width of 0.6s) seems to give an appropriate level of smoothing. This method is used in the remainder of this study.

Figure 3 Normalised accelerometer data for a two-minute accelerometer window and the corresponding IAE computed by the three methods.
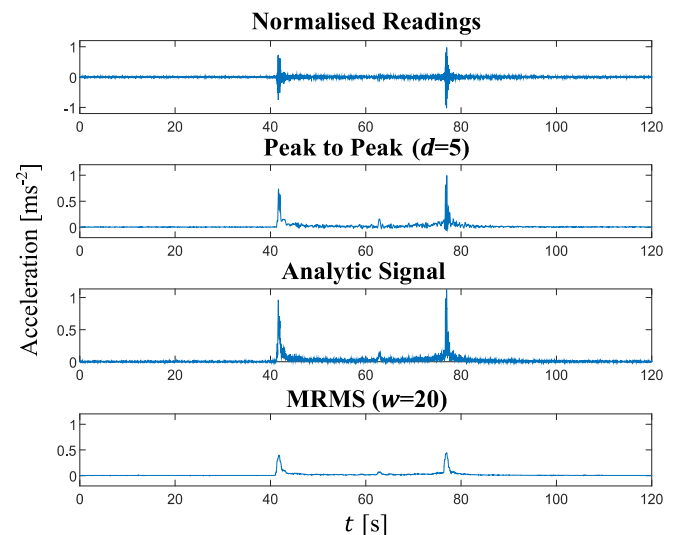


### 3.3 Labelled data set

The toll barrier data is inspected to find one-minute windows when one, two or three vehicles enter the bridge at one end, and no vehicles enter at the other end. Further, only windows

753

for which no vehicles enter at either end in the preceding and the following minute are considered. A two-minute window of IAE accelerometer data, corresponding to the selected minute and that which follows, is then extracted for analysis. Such quiet traffic conditions feature mostly at night-time, so the search was restricted to 10pm to 7am each night. In addition, this selection reduces the likelihood of structural responses to pedestrians or cyclists, whose ground truth numbers are not known.

In total, we built a training set of 30 windows and a validation set of 120 windows, as shown in Table 1. Here the counts from the Leigh Woods and Clifton toll barriers are denoted by $n_{\mathrm{LW}}$ and $n_{\mathrm{C}}$ respectively. The training windows are used to tune the method in Section 4 and the validation windows are used to test its accuracy in Section 5.

The method identifies the structural response to each vehicle, and thus automatically counts vehicles travelling across the bridge in each window. The method is developed to work both when the response is clear for each vehicle, as shown in Figures 2 and 3, and when the bridge is busier.

To summarise: each selected window is equipped with the count for each toll barrier $n_{\mathrm{LW}}$ and $n_{\mathrm{C}}$, and a smoothed IAE time series $\hat{a}_i$ for $i = 1, 2, \ldots, m$. For reference purposes, the time $T(t)$ (YYYYMMDD hh:mm:00) at which the window commences is also recorded.

Table 1 Labelled data set of 30 training windows and 120 validation windows. We collect 30 training windows with a single vehicle for training and 120 validation windows for one, two and three vehicles.

| $n_{\mathrm{LW}}$ | $n_{\mathrm{C}}$ | Training windows | Validation windows |
|---|---|---|---|
| 1 | 0 | 15 | 20 |
| 2 | 0 | 0 | 20 |
| 3 | 0 | 0 | 20 |
| 0 | 1 | 15 | 20 |
| 0 | 2 | 0 | 20 |
| 0 | 3 | 0 | 20 |

## 4. Method

The bridge will exhibit a damped oscillatory response to loading. This response is multi-modal and complex. The general idea is to seek the optimal fit of the IAE with a *simplified* closed-form function that captures the number of vehicles crossing the bridge. Section 4.1 describes the method and Section 4.2 fits it to the training set.

### 4.1 Derivation of the method

We try to fit a function to the smoothed IAE time series, $\hat{a}_i$ for $i = 1, 2, \ldots, m$. The IAE begins at time $t_0$, and the $i^{\mathrm{th}}$ reading

$\hat{a}_i$, is recorded at $t = t_0 + i\Delta t$. This provides a mapping from the discrete time index $i$ of the IAE, to continuous time $t$.

*4.1.1 Single vehicle:* We use the basis function $f(t; \boldsymbol{\theta})$ to fit the IAE of the bridge's structural response under the loading of a single vehicle. We assume that the function is smooth and is parameterised by $p$ parameters, $\boldsymbol{\theta} = [\theta_1, \theta_2, \ldots, \theta_p]$.

In addition, we include a scaling parameter $\gamma$ to account for different envelope amplitudes, and a shift parameter $\tau$ that locates the vehicle within the two-minute window. Thus we fit

$$g(t; \boldsymbol{\theta}, \gamma, \tau) = \gamma f(t - \tau, \boldsymbol{\theta}). \tag{4}$$

*4.1.2 Multiple vehicles:* Under a linear superposition assumption, we model the structural response to $n$ vehicles by the weighted sum of $n$ single-vehicle basis functions in the form

$$g_n(t) = \sum_{j=1}^{n} g_j(t; \boldsymbol{\theta}_j, \gamma_j, \tau_j) = \sum_{j=1}^{n} \gamma_j f(t - \tau_j; \boldsymbol{\theta}_j). \tag{5}$$

The function $g_n(t)$ is parameterised by $n(p + 2)$ parameters $\boldsymbol{\theta}_j = [\theta_{j_1}, \theta_{j_2}, \ldots, \theta_{j_p}]$, $\gamma_j$ and $\tau_j$ for $j = 1, 2 \ldots, n$.

The chosen estimate $n^*$ for the number of vehicles, is then found by the putative 'trial' number of vehicles $n$ that minimises the mean square error (MSE),

$$\mathrm{MSE} := \frac{1}{m} \sum_{i=1}^{m} \left( \hat{a}_i - g_n(t_o + i\Delta t) \right)^2, \tag{6}$$

between the fitted function and the IAE.

The parameters $\boldsymbol{\theta}_j$, $\gamma_j$ and $\tau_j$ are found in sequence. That is, the parameters that describe a single vehicle $\boldsymbol{\theta}_1$, $\gamma_1$ and $\tau_1$, are solved for and fixed. Then, the parameters that describe a second vehicle, $\boldsymbol{\theta}_2$, $\gamma_2$ and $\tau_2$, are solved for and fixed, and so on.

The procedure continues to add vehicles until the information criterion described in Section 4.1.3 tells it to stop. At this point, the procedure has arrived at its final estimate $n^*$.

*4.1.3 Akaike information criterion:* To avoid over-fitting, the Akaike information criterion (AIC) (Akaike, 2011) is used to penalise the goodness of fit of $g_n(t)$ by the number of used parameters $n(p + 2)$. The AIC is defined by

$$2n(p + 2) - 2\log \hat{L}, \tag{7}$$

where $\log \hat{L}$ is the maximum value of the log-likelihood function of the fit $g_n(t)$ given the IAE. The chosen estimate $n^*$ is selected by the function that minimises the AIC.

*4.1.4 Auto-correlation and the likelihood function:* Standard likelihood techniques make the assumption that the error terms between the fitted function and the IAE, $\varepsilon_i = \hat{a}_i - g_n(t_0 + i\Delta t)$, are identically and independently distributed (iid). This is not an appropriate assumption when dealing with time series data, as consecutive errors are likely to be correlated.

Instead, we assume that the errors follow an auto-regressive (AR) process. Each error is modelled as the weighted sum of the previous $q$ errors plus a noise term, which is assumed to be iid normally distributed with zero mean and variance $\sigma_\varepsilon^2$ in the form

$$\varepsilon_i = \sum_{k=0}^{q} \phi_k \varepsilon_{i-k} + N(0, \sigma_\varepsilon), \tag{8}$$

where $|\phi_k| < 1$ for $k = 0, 1, \ldots, q$. The parameter $q$ is known as the order of the AR model, which is denoted as AR($q$). The AR coefficients $\phi_k$ (weights) define how correlated each error $\varepsilon_i$ is with the $k^{\text{th}}$ error before ($\varepsilon_{i-k}$), known as the $k^{\text{th}}$ lag.

The AR coefficients are determined by the Partial Auto-Correlation Function (PACF) (Ramsey, 1974). In practice, software routines such as the parcorr function in Matlab can be employed to find the PACF. The choice of model order $q$ is described in Section 4.2.3

The $i^{\text{th}}$ error from the AR process is then defined by

$$\hat{\varepsilon}_i = \varepsilon_i - \sum_{k=1}^{q} \phi_k \varepsilon_{i-k}, \tag{9}$$

where $\hat{\varepsilon} \sim N(0, \sigma_\varepsilon)$.

Note the two-minute windows contain periods when there are no vehicles on the bridge, and therefore, there is no structural response, only background noise.

So that we only consider data above the noise floor, a threshold $\delta$ is set. We then consider only those $\hat{m} < m$ errors $\bar{\varepsilon}_i$ for which $\hat{a}_i > \delta$ for $i = 1, 2, \ldots, \hat{m}$. Here the threshold $\delta$ is set to 0.05ms$^{-2}$, by visual inspection.

The errors $\bar{\varepsilon}_i$ are also normally distributed with zero mean but with variance $\sigma_{\bar{\varepsilon}}^2$. A standard calculation gives the log-likelihood function for the fitted function $g_n(t)$ in the form

$$\log \hat{L} = -\frac{\hat{m}}{2} \log \left( \sum_{i=1}^{\hat{m}} \bar{\varepsilon}_i^2 \right) + c, \tag{10}$$

for constant $c$. Substituting Equation (10) into Equation (7) gives the quantity to be minimised. To summarise, in doing so we find the estimated number of vehicles that cross the bridge, from

$$n^* = \underset{n}{\text{argmin}} \left[ 2n(p+2) + \hat{m} \log \left( \sum_{i=1}^{\hat{m}} \bar{\varepsilon}_i^2 \right) \right]. \tag{11}$$

## 4.2 Training the method

To determine the basis functions and various parameters in the method, we consider its application to a training set in which only one vehicle is present in each two-minute window, see Table 1. The restriction to single-vehicle set-ups is in order to factor out additional problems if the linear superposition assumption turned out to be poor.

*4.2.1 Choice of basis functions:* We consider the basis function $f(t; \boldsymbol{\theta})$ that models the IAE time series of the bridge's structural response to a single vehicle entering the bridge at either end. A good basis function will give a small mean squared error in Equation (6) with $n = 1$.

Thus each selected IAE summarises the damped oscillatory response of the bridge to the dynamic loading caused by a single passing vehicle. The overall shape should be one with a rapid 'attack' and a gradual (probably exponential) decay. Inspired by probability theory, we considered the following three forms for the basis function, each of which has ($p = 2$) two parameters that control 'location' and 'dispersion' and thus the overall 'shape'.

**log-normal:**

$$f(t; \boldsymbol{\theta}) = \frac{1}{t\sqrt{2\pi}\sigma} \exp \left( -\frac{(\log(t) - \mu)^2}{2\sigma^2} \right),$$
$$\boldsymbol{\theta} = [\mu, \sigma] \tag{12}$$

**gamma:**

$$f(t; \boldsymbol{\theta}) = \frac{\beta^\alpha}{\Gamma(a)} t^{\alpha-1} \exp(-\beta t), \quad \boldsymbol{\theta} = [\alpha, \beta] \tag{13}$$

**Weibull:**

$$f(t; \boldsymbol{\theta}) = \frac{k}{\lambda} \left( \frac{t}{\lambda} \right)^{k-1} \exp \left( -\frac{t^k}{\lambda^k} \right), \quad \boldsymbol{\theta} = [k, \lambda] \tag{14}$$

The quality of each function is computed by calculating the optimal MSE for each of the 30 single-vehicle training windows. Table 2 shows the mean MSE (MMSE), i.e., the average across the 30 windows. The log-normal fit is best and so log-normal basis functions are used going forward.

*4.2.2 Examination of the log-normal parameters:* We now consider whether the log-normal parameters $\mu$ and $\sigma$ are similar for each vehicle travelling across the bridge. If they were, they might be fixed in the fitting process so that only $\gamma$ and $\tau$ need be determined for each vehicle.

Table 2 MMSE for the log-normal, gamma and Weibull fits to single-vehicle windows.

| Function | MMSE Leigh Woods [ms$^{-2}$] | MMSE Clifton [ms$^{-2}$] | MMSE Total [ms$^{-2}$] |
|---|---|---|---|
| log-normal | 0.0005 | 0.0012 | **0.0008** |
| gamma | 0.0006 | 0.0012 | 0.0009 |
| Weibull | 0.0017 | 0.0022 | 0.0019 |

For each of the training windows we fit a log-normal function by least squares to find the optimal parameters $\boldsymbol{\theta} = [\mu, \sigma]$, $\gamma$ and $\tau$. Figure 4 shows a scatter plot of the optimal $\mu$ and $\sigma$ parameters. There is unfortunately no consistent pattern and therefore the log-normal parameters must be found independently for each vehicle going forward.

Figure 4 Scatter plot of log-normal parameters fitted to 30 training windows.
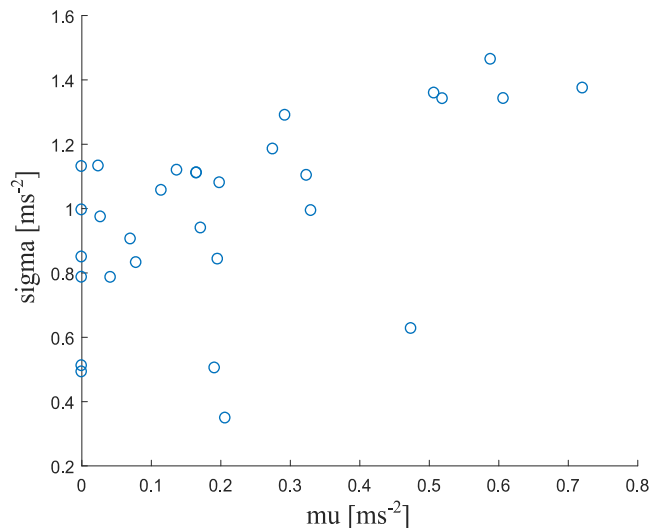


*4.2.3 Auto-regressive model order:* We now use the training windows to determine the order of the auto-regressive model $q$ described in Section 4.1.4. For each window we fit a log-normal function and compute the errors $\varepsilon_i$ and the first 30 weights $\phi_k$. We then calculate the mean weights $\bar{\phi}_k$ averaged over the training windows.

The model order is determined by the lag for which the mean weight falls within the 95% significance interval (Bollerslev, 1986). Figure 5 shows the mean weights for each of the tested lags. All of the significant autocorrelation is contained within the first eight lags. Therefore, in the likelihood calculation described in Section 4.1.4, the error terms $\hat{\varepsilon}$ are computed by an AR(8) process.

The AR model order $q$ is the only system parameter learnt prior to the vehicle counting procedure described in Section 4. In summary, for each window and for each putative number of vehicles, new log-normal means $\mu$ and standard deviations $\sigma$, scales $\gamma$ and shifts $\tau$ are found by least squares. In the likelihood calculation, new sets of AR coefficients $\phi_k$ for $k = 0, 1, \dots, 8$ are also computed.

Figure 5 Mean AR coefficients and the 95% confidence intervals.
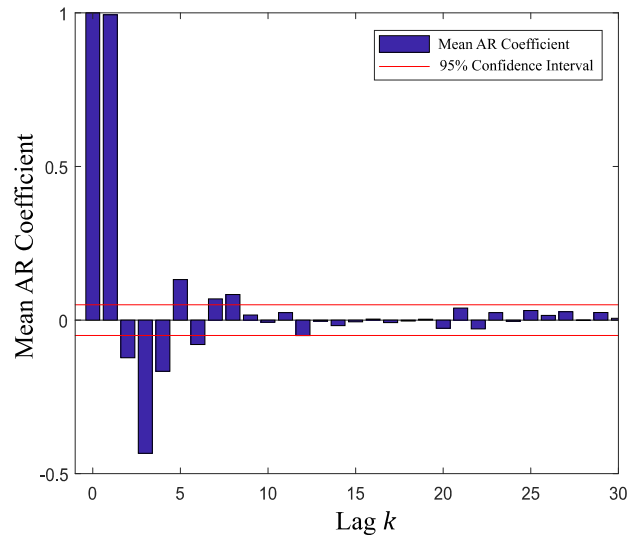


## 5. Results

We test the method by computing the estimated number $n^*$ of vehicles for 120 validation windows and comparing with their ground truth counts $n_{\text{LW}}$ and $n_{\text{C}}$, see Table 1. The results are shown in Table 3. Since the method does not identify the direction of travel, $n_{\text{LW}} + n_{\text{C}} = n^*$ corresponds to a correct prediction (although for any one validation window, only one of $n_{\text{LW}}$ or $n_{\text{C}}$ is non-zero). The method estimates the correct number of vehicles in 89 (74%) of the validation windows.

Table 3 Toll barrier counts and the corresponding estimated number of vehicles $n^*$, for 120 validation windows.

| $n_{\text{LW}}$ | $n_{\text{C}}$ | $n^* = 1$ | $n^* = 2$ | $n^* = 3$ | $n^* = 4$ |
|---|---|---|---|---|---|
| 1 | 0 | 19 | 1 | 0 | 0 |
| 2 | 0 | 0 | 18 | 2 | 0 |
| 3 | 0 | 0 | 3 | 15 | 2 |
| 0 | 1 | 13 | 4 | 3 | 0 |
| 0 | 2 | 0 | 12 | 6 | 2 |
| 0 | 3 | 0 | 2 | 12 | 6 |

Figure 6 illustrates the method at work for two exemplar windows. In each example, the MSE decreases as the putative number of 'trial' vehicles increases and the function $g_n(t)$ becomes more complex. However, the AIC penalises the

Figure 6 Estimation method in operation for two exemplar windows. (a) $n_{LW} = 2, n_C = 0$ (b) $n_{LW} = 0, n_C = 1$. Left-hand panels (i) show the superposition of one, two, three and four basis functions fitted to the IAEs. Right-hand panels (ii) show the corresponding MSE and AIC for the fits.

(a)(i)

(a)(ii)

(b)(i)

(b)(ii)
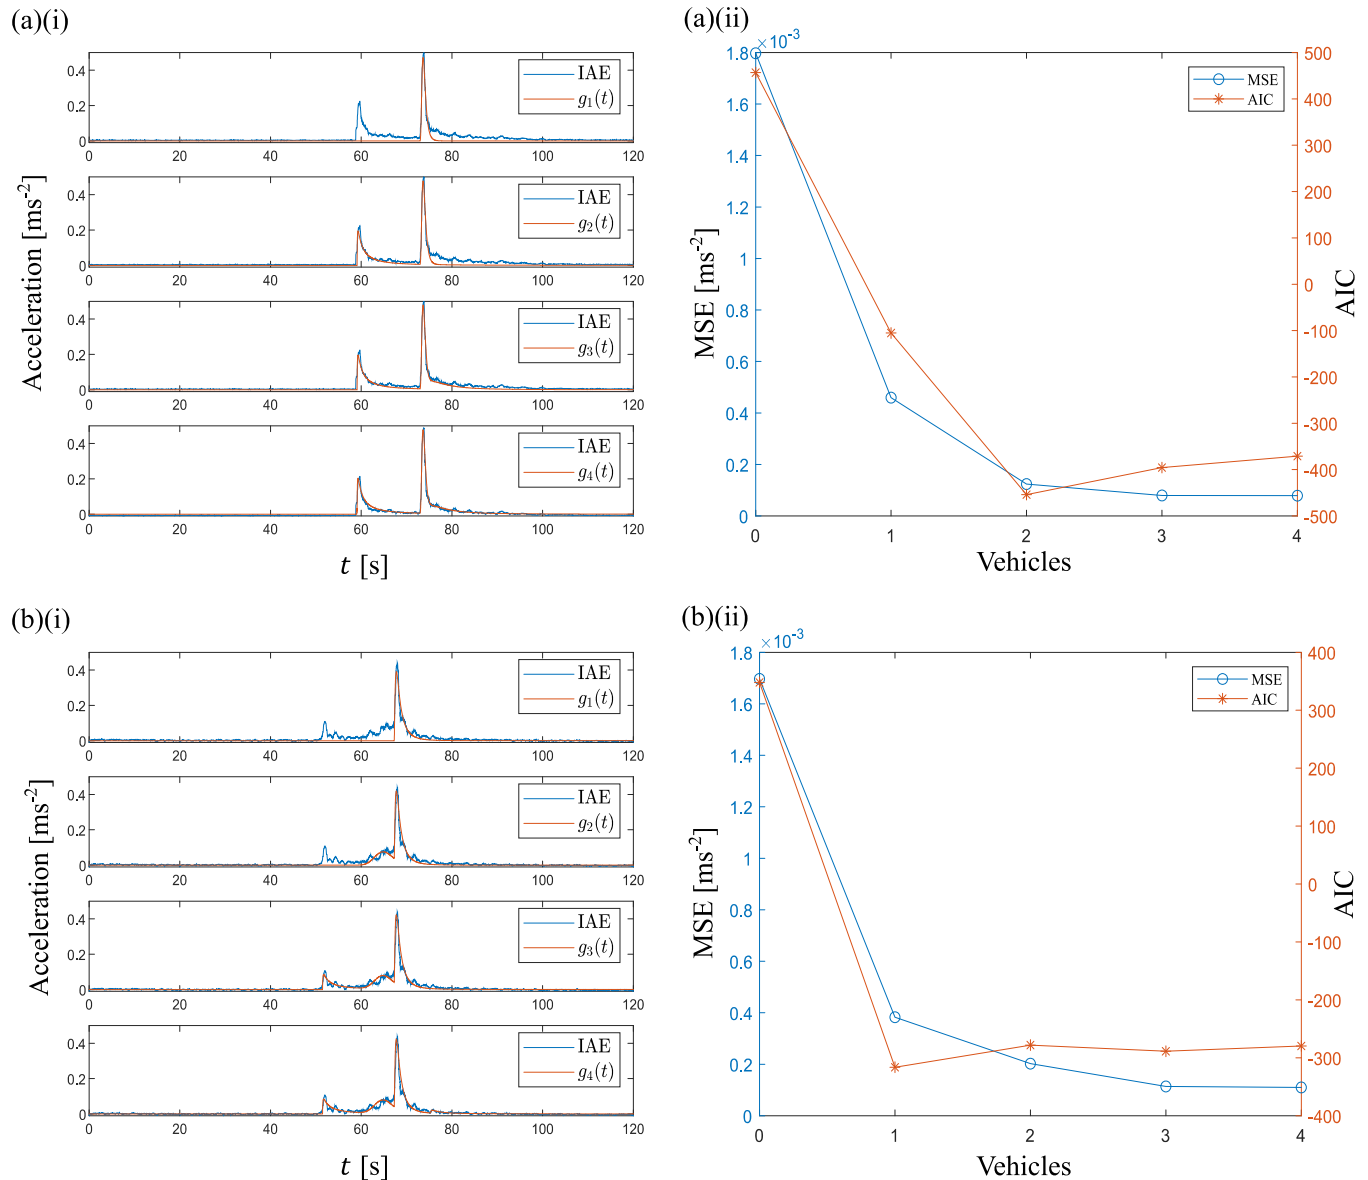
complexity of $g_n(t)$, and the number of vehicles is found correctly, by its minimisation.

## 6. Discussion

The overall prediction accuracy of 74% is somewhat disappointing. However, when $n_C = 0$ (vehicles only enter at the Leigh Woods end), we achieve an accuracy of 87% and we have the beginnings of a plausible operational method. In addition, for these windows, there is no evidence of bias towards over or under-counting.

In contrast, the fit to vehicles entering at the Clifton end is poor, and there is apparently a bias towards over-counting. This may

be explained by the exemplar window shown in Figure 7. A single log-normal basis function is a poor fit to the structural response in this case, because the vehicle causes two peaks in the structural response, one when it enters the bridge at the Clifton end, and a second when it reaches the accelerometer at the Leigh Woods end. The estimation method attempts to fit the two peaks with two basis functions, and so counts two vehicles when only one is present. This problem does not arise for vehicles entering at the Leigh Woods end, because the entry peak is almost simultaneous with that when the accelerometer is reached. Note Table 2, which concerns the training set, has already indicated that unimodal basis functions fit the Leigh Woods vehicles better than the Clifton ones.

Further work should consider different basis functions for the two directions of travel, which would also allow the counts for the two directions of travel to be separated. Alternatively, we might place the accelerometer in a position which is indifferent to the direction of travel, e.g., midway along the span. Note also that this paper has only used the data from a single accelerometer, as the aim was to demonstrate the beginnings of a plausible operational method that might be applied to existing SHM set-ups, with a very limited number of sensors. However, clearly a range of improvements might be possible if the data from multiple accelerometers were fused.

Figure 7 An example where $n_C = 1$, $(n_{LW} = 0)$ and $n^* = 2$. When the vehicle enters at the Clifton end, there is an additional peak in advance of that when it reaches the Leigh Woods end detector.
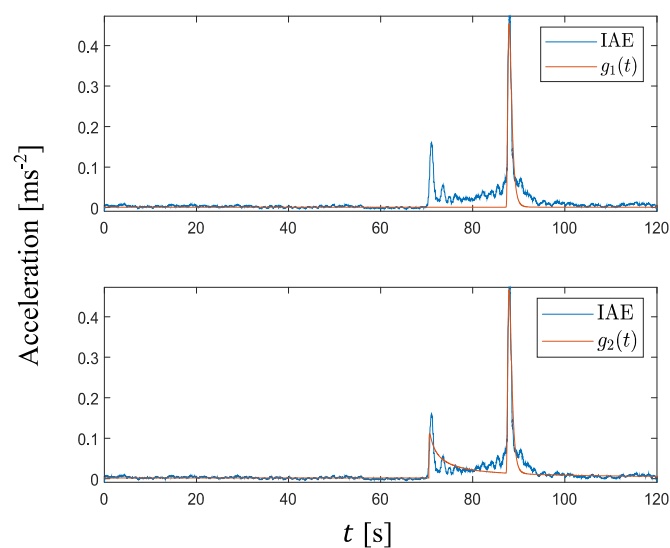


## 7. Conclusion

In this paper we have developed a method that estimates the number of vehicles that cross a bridge from its structural response to their dynamic loading. We have applied the method to scalar accelerometer data from the Clifton Suspension Bridge (CSB) in Bristol, UK. A validation exercise applied to a set of selected two-minute windows of data counted the correct number of vehicles in 74% of the set. However, due to various features of the set-up, which was designed for structural health monitoring (SHM), not for counting vehicles, the method performs quite differently in the two directions of travel – 61% and 87% accuracy respectively. The higher figure suggests that with some refinements to the method and sensor set-up, we have the beginnings of a deployable operational system. Note that for many practical traffic management and appraisal purposes, the key is not that vehicles are counted precisely, but rather that an offline estimate of the count is unbiased over quite long time periods – this requirement is relatively tame and looks to be within reach.

There are several shortcomings in our current work and consequently opportunities for future research. Firstly, as currently presented, the method does not distinguish between the directions of travel. However, this appears addressable by considering the data from multiple sensors and distinct basis functions for each direction. Secondly, the validation set currently used is not representative of typical operational conditions, because it does not mix vehicles with different directions of travel, and the total vehicle counts are quite low, compared to busier times of the day. Future work should investigate to what degree the method can identify individual vehicles when there are many different overlapping structural responses, perhaps including those from cyclists and pedestrians too, and indeed whether linear superposition holds in congested traffic flow regimes. Finally, the method as presented works because the CSB is much more compliant than modern concrete bridges, where an examination of higher frequency information, rather than the amplitude envelopes used here, is likely to be the way forward.

## 8. References

Akaike, H., 2011. Akaike's information criterion. In: M. Lovric, ed. *International Encyclopedia of Statistical Science.* Berlin: Springer Berlin Heidelberg, pp. 25-25.

Albu, M. & Heydt, G. T., 2003. On the use of RMS values in power quality assessment. *IEEE Transactions on Power Delivery,* Volume 18, pp. 1586-1587.

Basharat, A., Catbas, N. & Shah, M., 2005. A framework for intelligent sensor network with video camera for structural health monitoring of bridges. *Third IEEE International Conference on Pervasive Computing and Communications Workshops,* pp. 385-389.

Biswas, A. & Si, B., 2011. Depth persistence of the spatial pattern of soil water storage in a hummocky landscape. *Elements,* Volume 29, pp. 19-23.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics,* Volume 31, pp. 307-327.

Chan, T. H. T. et al., 2006. Fiber Bragg grating sensors for structural health monitoring of Tsing Ma bridge: Background and experimental observation. *Engineering Structures,* Volume 28, pp. 648-659.

Gunner, S. et al., 2017. Rapid deployment of a WSN on the Clifton Suspension Bridge, UK. *Proceedings of the Institution of Civil Engineers - Smart Infrastructure and Construction,* Volume 170, pp. 59-71.

Hayes, J. G. & Halliday, J., 1974. The least-squares fitting of cubic spline surfaces to general data sets. *IMA Journal of Applied Mathematics,* Volume 14, pp. 89-103.

Macdonald, J. H. G., 2008. Pedestrian-induced vibrations of the Clifton Suspension Bridge, UK. *Proceedings of the ICE - Bridge Engineering,* Volume 161(2), pp. 69-77.

Oh, B. K. et al., 2015. Vision-based system identification technique for building structures using a motion capture system. *Journal of Sound and Vibration,* Volume 356, pp. 72-85.

Picinbono, B., 1997. On instantaneous amplitude and phase of signals. *IEEE Transactions on Signal Processing,* Volume 45, pp. 552-560.

Ramsey, F. L., 1974. Characterization of the partial autocorrelation function. *Ann. Statist.,* Volume 2, pp. 1296-1301.

Reilly, A., Frazer, G. & Boashash, B., 1994. Analytic signal generation-tips and traps. *IEEE Transactions on Signal Processing,* Volume 42, pp. 3241-3245.

Yeung, W. T. & Smith, J. W., 2005. Damage detection in bridges using neural networks for pattern recognition of vibration signatures. *Engineering Structures,* Volume 27, pp. 685-698.