

Characterisation of pavement profile heights using accelerometer readings and a combinatorial optimisation technique

N.K. Harris, A. Gonzalez¹, E.J. OBrien & P. McGetrick

School of Architecture, Landscape & Civil Engineering
University College Dublin
Newstead
Belfield
Dublin 4

¹ Corresponding author
Phone: +353-1-7163219
Fax: +353-1-7163297
e-mail: arturo.gonzalez@ucd.ie

ABSTRACT

Pavement surface profiles induce dynamic ride responses in vehicles which can potentially be used to classify road surface roughness. A novel method is proposed for the characterisation of pavement roughness through an analysis of vehicle accelerations. A combinatorial optimisation technique is applied to the determination of pavement profile heights based on measured accelerations at and above the vehicle axle. Such an approach, using low-cost inertial sensors, would provide an inexpensive alternative to the costly laser based profile measurement vehicles. The concept is numerically validated using a half-car roll dynamic model to infer measurements of road profiles in both the left and right wheel paths.

Keywords: road profiles; parameter identification; combinatorial optimisation

1. INTRODUCTION

Road profiles are an intrinsic part of the interaction between a pavement and/or bridge and a heavy vehicle suspension. As such, it is generally accepted that the maintenance of smooth road profiles for highways and bridges is of major importance in minimising dynamic tyre forces, promoting long pavement life spans and ensuring that bridge loads are small [1,2,3]. This paper outlines a novel approach for the characterisation of road profiles using low cost accelerometer measurements.

Currently, there exist several methods for the measurement of road surface profiles. Static methods, such as the use of rod and level equipment or specialised ‘dipstick’ walking profilometers, are slow and time consuming. The use of dynamic methods, such as inertial profilometers which measure single (and often multiple) profile tracks at typical highway velocities, allows for large scale and continuous pavement measurement regimes. The typical inertial profilometer consists of a vehicle equipped with a height sensing device, such as a laser, which measures pavement elevations at regular intervals [4,5]. Accelerometer(s) mounted on the vehicle allow the effects of vehicle dynamics to be removed from the elevation measurements, resulting in a measured profile. This method provides an accurate, high resolution measurement of road profile, though the associated costs of laser-based technology are a disadvantage. Imine et al [6] describe a technique for road profile input estimation based on the measured dynamic response of an instrumented vehicle. The method proposed uses a full car sprung mass model to estimate profile heights using as input the wheel vertical accelerations and vertical

displacement and rotation of the vehicle body. The authors compare the results obtained to those of an inertial profilometer and to that of the ‘longitudinal profile analyzer’; a profile measuring device designed at Laboratoire Central des Ponts et Chaussées (LCPC). Fair agreement is found for the proposed method, with noticeable local discrepancies in the estimated profile.

González et al [7] propose an easy-to-implement road classification method based on the relationship between the power spectral densities of vehicle accelerations and road profile via a transfer function. This frequency-domain method generally classifies the road into the appropriate class, but it is unable to calculate the height of the road irregularities at each point in time.

In order to infer a desired series of road profile elevations from measurements of vehicle response to them, the new method proposed herein uses of a combinatorial optimisation technique, known as the cross-entropy method [8]. Initially, the parameters of a representative vehicle model are determined by examining the vertical acceleration response of the vehicle to a known excitation (e.g. an obstacle such as bump, or a road profile which is already known). Using this calibrated vehicle, unknown road profiles may then be characterised by optimising to find the road profile which replicates the measured response. The proposed method is tested using simulated acceleration measurements, taken from the unsprung and sprung masses for a number of road profiles.

2. CROSS-ENTROPY METHOD FOR COMBINATORIAL OPTIMISATION

The inference of unknown model parameters is a complex and computationally intensive problem, usually requiring the exploration of large, multidimensional solution spaces in order to yield an answer. This process, known as combinatorial optimisation, may be informally defined as a problem where the set of feasible solutions is discrete and the goal is to determine the best possible solution. The preferred optimisation method used for this study is the Cross-Entropy (CE) method which is an iterative approach that employs Monte Carlo simulation to arrive at an optimal solution. The CE method is chosen due to its relative ease of implementation and its ability to avoid becoming stuck in local optima [8,9].

The CE method is an iterative process, which involves two distinct phases:

- 1) Generation of a sample of random data (e.g. in this case, vehicle model parameters or pavement elevations) according to a specified random mechanism.
- 2) Updating of the parameters of the random mechanism in order to produce an improved sample in the next generation.

More formally, the process, adapted from De Boer et al [8], is described as follows:

Consider a finite set of states, X , which, for example, may contain all possible values for the parameters of some vehicle ride model. Let S then be a real-valued function on X

which describes some measurable output, such as vertical accelerations. There must then be an array of optimal states, χ^* , that maximises S , i.e.

$$S(\chi^*) = \max_{\chi \in X} S(\chi) \quad (1)$$

Eq. (1) is a general definition of an optimisation problem. When χ consists of more than one parameter, it becomes a combinatorial optimisation problem. In order to arrive at an estimation of χ^* , a set of N_χ iid (independent and identically distributed) random samples of χ are generated from the probability distribution characterised by the set of parameters, ν (e.g. mean, standard deviation):

$$x_1, x_2, \dots, x_N \sim f(\chi | \nu) \quad (2)$$

where $f(\chi | \nu)$ is some probability distribution (e.g. Gaussian, lognormal, etc.). Usually, the distribution is assumed to be Gaussian and thus ν contains the means and variances of the distribution. Because at the beginning of the optimisation process, little is known about the parameters, one may select arbitrary mean values and very large variances to reflect the lack of knowledge about possible values for χ . For each randomly generated $\chi_{1,2,\dots}$ the performance function, $S(\chi)$, is evaluated. From this, the sample $(1-\rho)$ quantile is determined, i.e. if $\rho = 0.9$, then the top 10 % of answers are taken. An updated estimate, $\hat{\nu}$, of the probability distribution parameters is then obtained by taking the mean and variance of the samples from the $(1-\rho)$ quantile. This process is repeated until a

stopping criterion is reached, such as a finite number of generations or convergence of estimates.

Fig. 1 outlines the general CE algorithm as implemented for this study. The stopping criterion for the algorithm is simply based on the number of generations. This was found, empirically, to be sufficient. After half of the generations are completed, the variance parameters in $\hat{\nu}$ are widened in order to reduce the chances of becoming stuck in a local maximum of the solution space.

[FIG. 1 HERE]

Fig. 2 shows the operation of the CE method for a simple mathematical function with two maxima. The objective is to find the value of y which maximises the function $f(y)$ (the solid line in the graph). The dashed line shows the probability distribution used to draw samples, χ , for each generation, shown as circles. As the optimisation process progresses, the sampling distribution is updated by taking the mean and variance of the top 10 % of values which maximise $f(y)$. The algorithm swiftly converges to the maximum value of $y=2$.

[FIG. 2 HERE]

3. CHARACTERISATION OF VEHICLE MODEL PARAMETERS

Initially, it is necessary to determine the unknown vehicle model parameters based on observed measurements of vehicle response due to a known disturbance, e.g. a bump or a known road profile. This may be subsequently used to specify a reference vehicle model for the characterisation of unknown pavement profile heights. A half-car roll model (Fig. 3) is used for this purpose since it will be shown later that it is possible to infer two parallel profiles in the left and right wheel paths with this model.

[FIG. 3 HERE]

The half-car roll model has four DOFs to account for the axle hop, sprung mass heave and sprung mass rolling motions induced by different road profiles underneath the left and right wheel paths. By assuming negligible tyre damping [10], there are up to 11 unknown parameters associated with the model; the vehicle sprung mass and roll inertia, the unsprung mass, suspension stiffness, damping coefficient and tyre stiffness for both the left and right sides and, finally, the track width of the model. By assuming symmetry (geometrically and of suspension components) and deeming the track width, $2T$, to be easily measurable, the model is thus reduced to six unknowns: m_s , I_s , m_u , k , k_t , and c_b .

The half-car equations of motion for the sprung mass heave, y_s , sprung mass roll, ϕ_s and axle hop motions, y_{u1} , y_{u2} , respectively are given by:

$$m_s \ddot{y}_s + F_{s1} + F_{s2} = 0 \quad (3)$$

$$I_s \ddot{\phi}_s + TF_{s1} - TF_{s2} = 0 \quad (4)$$

$$m_u \ddot{y}_{ui} + F_{si} - F_{ri} = 0; \quad i = 1, 2 \quad (5)$$

Assuming a small roll angle, the quantity F_{si} , which denotes the suspension force generated above wheel i , can be expressed by Eq. (6), while the corresponding tyre force at wheel i , F_{ri} , is given by Eq. (7).

$$F_{si} = k \left(y_s - (-1)^i T \phi_s - y_{ui} \right) + c_b \left(\dot{y}_s - (-1)^i T \dot{\phi}_s - \dot{y}_{ui} \right); \quad i = 1, 2 \quad (6)$$

$$F_{ri} = k_t (y_{ui} - r_i) \geq -P_i; \quad i = 1, 2 \quad (7)$$

Eq. (7) imposes a lower bound on dynamic tyre force, allowing the model to lose contact with the road if the upwards dynamic force exceeds the static weight, P_i , of the vehicle at wheel i , given by:

$$P_i = g(m_u + 0.5m_s); \quad i = 1, 2 \quad (8)$$

where g is the acceleration due to gravity. Table 1 lists the vehicle model parameters used for the simulations. These represent the ‘true’ answer sought through the CE optimisation.

[TABLE 1 HERE]

3.1 Objective function for optimisation

By the CE optimisation process given in Section 2, a performance function is defined to evaluate the suitability of solutions proposed by the algorithm. This is taken to be a relatively simple least squares minimisation of the difference, over N_j time steps, between the observed data and the simulated outputs (using the proposed solutions) of the vehicle model. It is represented by the function $S(\chi)$ in Eq. (9) that needs to be maximized.

$$S(\chi) = \left(\sum_{j=1}^{N_j} (A_{\text{obs}}(t_j) - A_{\text{sim}}(\chi, t_j))^2 \right)^{-1} \quad (9)$$

where t_j is the j^{th} time step. The vehicle response functions, $A_{\text{obs}}(t)$ and $A_{\text{sim}}(\chi, t)$ are specified based on the combination of acceleration measurements taken from the vehicle.

3.2 Road profile inputs

The characterisation of the vehicle model parameters is investigated for both an artificially generated road profile and for a simple bump input. These are assumed to be known inputs to the simulation model. The road profile used is 100 m in length and is of ISO class A roughness (good quality highway) [11]. A sinusoidal bump of amplitude 20 mm is also used, with bump wavelengths of 0.25, 0.5, 1.0, 1.5 and 2.0 m. In each case, the vehicle velocity is constant at 80 km hr⁻¹.

3.3 Sources of ‘observed’ data

This study considers the use of acceleration measurements as the reference sources of observed data for the optimisation. For the half-car model, it is assumed that four measurement sources are available – the accelerations at each wheel and the sprung mass accelerations above each suspension, \ddot{y}_{s1} and \ddot{y}_{s2} , given by:

$$\ddot{y}_{si} = \ddot{y}_s - (-1)^i T \ddot{\phi}_s; \quad i = 1,2 \quad (10)$$

The vehicle response function, introduced in Eq. (9), is thus written as:

$$A(t) = \eta_1 \ddot{y}_{u1}(t) + \eta_2 \ddot{y}_{u2}(t) + \eta_3 \ddot{y}_{s1}(t) + \eta_4 \ddot{y}_{s2}(t) \quad (11)$$

where η_1 to η_4 are weighting parameters which denote the relative importance each acceleration source has in the response function. To assign equal importance to each measurement source in the function:

$$\begin{aligned} \eta_1 &= \text{RMS}(\ddot{y}_{u2}(t)) \times \text{RMS}(\ddot{y}_{s1}(t)) \times \text{RMS}(\ddot{y}_{s2}(t)) \\ \eta_2 &= \text{RMS}(\ddot{y}_{u1}(t)) \times \text{RMS}(\ddot{y}_{s1}(t)) \times \text{RMS}(\ddot{y}_{s2}(t)) \\ \eta_3 &= \text{RMS}(\ddot{y}_{u1}(t)) \times \text{RMS}(\ddot{y}_{u2}(t)) \times \text{RMS}(\ddot{y}_{s2}(t)) \\ \eta_4 &= \text{RMS}(\ddot{y}_{u1}(t)) \times \text{RMS}(\ddot{y}_{u2}(t)) \times \text{RMS}(\ddot{y}_{s1}(t)) \end{aligned} \quad (12)$$

where RMS denotes the root mean square of the specified acceleration response. It is also important to note that the use of accelerations as target measurements means that information about the absolute values of the unknown parameters cannot be obtained. Rather, the algorithm must determine the relative ratios of each unknown with respect to

the others. The optimisations discussed herein are initiated by setting the unknown sprung mass parameter for both vehicle models to be fixed at an arbitrary value (say, 100 kg) and determining the other unknowns relative to this. The success of the optimisation may then be judged on the accuracy of these ratios.

3.4 Results of optimisation

The optimisation is run for the vehicle model at 80 km hr^{-1} , using the class A road profile as a known input and the measured accelerations as the known outputs for the problem. The CE algorithm uses a population size, N_x , of 250 samples for two stages of 20 generations. As discussed, the sprung mass unknown is held constant at an arbitrary value and thus the number of unknowns is reduced to five. These unknown parameters are sought using the CE optimisation process. The optimisation is tested to obtain the best parameters for 10 runs of the algorithm using the class A road (with different, but correlated left and right wheel paths) and the five sinusoidal bumps of varying wavelengths. In order to induce roll in the model, the sinusoidal bumps are only applied to the left wheel. The results, shown in Fig. 4, indicate the least error in estimation of parameters occurs for the pair of road profiles, whilst the bumps do not seem to be effective, particularly for determining roll inertia, I_s .

[FIG. 4 HERE]

Fig. 5 shows the unsprung mass accelerations for the vehicle model using the optimised parameters determined with the class A road as the known input, indicating an excellent

match to the observed data. To assign a numerical value to the quality of fit, the normalised RMS error parameter (NRE) is used [11]. This is given by the RMS error between the observed and predicted vehicle performance functions for the half-car (Eq.(11)), as a percentage of the RMS of the corresponding observed vehicle performance function:

$$\text{NRE} = \frac{\text{RMS}(A_{\text{obs}}(t) - A_{\text{sim}}(\chi^*, t))}{\text{RMS}(A_{\text{obs}}(t))} \times 100 \quad (13)$$

The NRE for the best fit vehicle response function is determined to be 9.3 %. Given that the NRE appears to be quite a sensitive measure of accuracy, these results are considered to be satisfactory, particularly since it is clear from the graph that good fits are obtained.

[FIG. 5 HERE]

4. CHARACTERISATION OF PAVEMENT PROFILE HEIGHTS

The previous section describes the determination of unknown vehicle model parameters based upon the known excitations due to the road profile. As such, the problem consists of determining the parameters of the half-car model using known inputs (road profile(s)) and known outputs (accelerations). This section discusses finding the problem of unknown road profile heights in the same manner. Here, the problem to be addressed involves optimising to find the unknown inputs to a fully characterised process model

(known vehicle parameters, etc.) which cause a set of known acceleration responses. Again, the CE method is used to determine the optimal solution set.

4.1 Optimisation process for unknown profile heights

The process employed for the characterisation of a set of unknown profile heights is as follows:

1. The unknown road profile is idealised as an optimisation problem with N_h unknowns, where $N_h = L/\Delta$ is the number of unknown changes in height between each interval for a length, L , of road with spacing, Δ . For this study, Δ is taken as 0.1 m, which corresponds to a spatial frequency of 10 cycles/m (upper bound in the ISO definition of power spectral density functions for the generation of artificial road profiles [11]). If two profile paths are being determined, the number of unknowns is doubled.
2. The optimisation problem is split into a number of phases, in which a smaller ‘window’ of the road profile is determined before proceeding to the next phase. This greatly reduces the dimensionality of the optimisation problem. A 1.0 m window is used, meaning that only 20 unknowns for the half-car are considered at any one time.
3. The CE optimisation process is used to determine the optimal set of height changes within the window which maximises the performance function, $S(\chi)$. Initially, all height changes are assumed to have a mean of 0 mm and a standard deviation of 1 mm.

4. When the profile heights in the window have been optimised, the algorithm proceeds to the next phase. The first half of the window is retained (i.e. 0.5 m of road heights), thus ‘building’ the unknown profile as the phases progress.
5. The heights in the second half of the window are used as the starting elevations for the next phase (i.e. locations 6 –10). The starting elevations for locations 11 – 15 are the same as for location 10. For each unknown elevation in the next phase, the initial standard deviation is increased on a sliding scale from 0.1 mm at location 6 to 1.0 mm at location 15. This is to reflect the relative uncertainty in estimates further along the window being analysed.

The process is repeated until all N_h elevations have been characterised. These form the estimated road profile, r_{est} , which may subsequently be post-processed to determine ride statistics such as the IRI. IRI is a standardized roughness measurement of the longitudinal profile of a wheel track, it has units of slope and it is calculated as a filtered ratio of a standard vehicle’s accumulated suspension motion (i.e. in mm) divided by the distance travelled by the vehicle during the measurement (i.e. in km) [5].

5. NUMERICAL VALIDATION OF ALGORITHM

The algorithm is numerically validated using several artificially generated road profiles of varying roughness. Each of the profiles *ArtificialRoad*1 – AR5, shown in Fig. 6, are 100 m in length, range from class A to class C roughness according to ISO standards [11] and have IRI ratings ranging from 1.07 – 8.58 m km⁻¹. In each case shown, the profile for the right wheel path is offset from the left path for clarity. Table 2 lists, for

each profile, the reference spatial frequency, $G_d(n_0)$, used to generate the profiles, the corresponding IRI values and the correlation, ρ_r , between the left and right profiles. All post-processing work (such as filtering, calculation of IRI and analysis of spectral densities) carried out on the road profiles in this study is done using ProVAL (**Profile Viewing and AnaLysis**; [12]). This is a freely available dedicated software package for such analyses, developed through the US Federal Highway Administration (FHWA) and the Long Term Pavement Performance Programme (LTPP).

[FIG. 6 HERE]

[TABLE 2 HERE]

For each of the road profiles, the half car vehicle model is run at 80 km hr^{-1} and the acceleration responses are taken for use as the inputs to the profile inference algorithm. Finally, the assumption is made here that the vehicle parameters can be fully specified using the approach described in Section 3 or an alternative procedure – thus the actual parameters (given in Table 1) are used in the validation. The approach proposed here to calculate the road profile using CE is also applicable to other vehicle models or configurations once the model parameters have been defined accurately.

The computational time required for the algorithm to estimate the profile of both wheel paths for one of the profiles shown in Fig. 6 was about 5 hours with a 3 GHz processor, 6 MB cache and 3072 MB SDRAM running on Matlab. The results of the road profile

inferences for AR2 (class A) and AR4 (class B/C) are shown in Figs. 7(a) and (b) respectively. The error in the estimation at each location is plotted in a secondary vertical axis. For clarity, only the last 25 m of the estimation are represented. This last quarter of the profile is where errors are more obvious although still small compared to the actual road irregularities. It can be seen that the algorithm has determined the road profile in the left and right wheel paths, capturing both the short and long wavelength components to a high degree of accuracy. It has been noted that in some cases, there is a gradual drifting effect to the estimated road profile, increasing with distance from the origin. This is an unavoidable effect associated with errors in the estimation of the target accelerations. Any differences in the simulated and observed accelerations will manifest as larger quadratic errors in the estimated profile because they are double-integrated with respect to time.

The corresponding actual and estimated power spectral densities (PSDs) of the profile heights for the left wheel path are shown in Fig. 8. Again, good agreement is obtained between actual and estimated profiles. Similar levels of accuracy are obtained for the profiles of the right wheel path. The main differences in the respective PSDs occur for the far left and right ends of the spectrum. There are some errors in estimating the high frequency, short wavelength components of the road profile (wavenumbers corresponding to excitation inputs greater than 50 Hz) and also some small errors for low frequency, long wavelengths, though this is less clear due to the logarithmic plot. These low frequency errors are a manifestation of a drifting effect.

[FIG. 7 HERE]

[FIG. 8 HERE]

It is also found that in order to initialise the algorithm correctly, it is necessary to specify the initial relative difference in the heights of both profile paths. This is because the algorithm finds all profile elevations relative to an initial profile height of 0 m. If one profile path is initially higher than the other then this information must be included when employing the algorithm. It is assumed that this initial relative height difference can be easily measured using an inclinometer or similar device and as such, is considered a known input.

5.1 IRI ratings for characterised profiles

The IRIs of the estimated profiles are calculated and compared to the true values from the original profiles. The results are summarised in Table 3. Each characterised road profile has an IRI value within $\pm 2\%$ of those for the true profiles. This indicates that in theory, the method proposed has the potential to measure road roughness to a class I standard [6]. Further, the accuracy of the method does not appear to be influenced by the level of roughness present, with similar accuracies obtained for profiles AR1 through to AR5.

[TABLE 3 HERE]

Despite this, the accuracy of the method would almost certainly be lower outside the realms of a numerical validation. The potential errors associated with modelling approaches (e.g. the assumption of linearity or mis-estimation of model parameters) or cumulative drift effect from double integrating accelerometer readings could lead to a reduction in accuracy in practice. Noise is other source of inaccuracy that will introduce deviations in the estimated profile.

5.2 Influence of noise on accuracy of the algorithm

The input half car acceleration signals for profiles AR2 and AR4 are corrupted with noise using an additive noise model with a signal to noise ratio of 20 (i.e., noise is randomly added to the true accelerations by sampling a normal distribution of zero mean and standard deviation equal to the standard deviation of the true acceleration data divided by 20). Fig. 9 shows 10-point moving average trends for the differences between the estimated and the actual road profiles when using noise-free or corrupted accelerations. Fig. 10 shows the actual, estimated and the noise-corrupted PSDs of road profile heights for the left wheel path. The figures show similar results to those for the uncorrupted data, except that the gradual drifting effect which was seen for the original estimated road profiles increases slightly with noise and higher road roughness. This drifting effect will become more important for higher levels of noise. The main differences in the respective PSDs occur for the far left and right ends of the spectrum as for the original estimated profiles.

[FIG. 9 HERE]

[FIG. 10 HERE]

Table 4 shows the estimated IRI for the level of noise being tested with no significant loss of accuracy with respect to the noise-free data.

[TABLE 4 HERE]

6. SUMMARY

A novel method for the characterisation of road surface profiles using measurements of vehicle acceleration response has been described. The method proposes the use of a combinatorial optimisation technique to determine the road input which causes a set of observed responses in a known vehicle model. The process is described for the characterisation of two parallel profiles, using a half-car roll model. Initially, the specification of a vehicle ride model for use in the algorithm is discussed. Using a known input, it is shown that the parameters of a half-car model can be found. It is also noted that the best results are obtained when the known excitation input has a wide enough spectral content to excite the vehicle model sufficiently.

The pavement characterisation algorithm is numerically validated for several artificial road profiles. The estimated road heights are found to provide a very good fit to the true

profiles for both models considered, with the IRI values consistently within $\pm 2\%$ of the reference values. The errors in the estimated profiles are generally confined to frequencies which have little influence on the vehicle model responses, such as short wavelength, high frequency errors and long wavelength errors due to drift. The addition of low levels of noise to the simulated accelerations has not affected the accuracy of the profile estimation significantly.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the support received from the 7th European Framework project *ASSET* (Advanced Safety and Driver Support in Efficient Road Transport) and the *Embark* scholarship, administered by the Irish Research Council of Science, Engineering and Technology (IRCSET).

REFERENCES

- [1] Dynamic interaction of heavy vehicles with roads and bridges. Technical report, OECD, Ottawa, Canada. DIVINE Concluding Conference, 1997.
- [2] T.D. Gillespie, S.M. Karamihas, D. Cebon, M.W. Sayers, M.A. Nasim, W. Hansen, N. Ehsan, Effects of heavy vehicle characteristics on p response and performance. University of Michigan Transportation Research Institute, UMTRI-92-2, 1992.

- [3] M.F. Green, D. Cebon, Dynamic responses of highway bridges to heavy vehicle loads, theory and experimental validation. *Journal of Sound and Vibration* 170 (1994) 51- 78.
- [4] M.W. Sayers, S.M. Karamihas, Interpretation of road profile roughness data. University of Michigan Transportation Research Institute, UMTRI-96-19, 1996.
- [5] M.W. Sayers, S.M. Karamihas, The little book of profiling. University of Michigan Transportation Research Institute, UMTRI-96-19, 1998.
- [6] H. Imine, Y. Delanne, N.K. M'Sirdi, Road profile input estimation in vehicle dynamics simulation. *Vehicle System Dynamics* 44 (2006) 285-303.
- [7] A. González, E.J. OBrien, Y.-Y. Li, K. Cashell, The use of vehicle acceleration measurements to estimate road roughness, *Vehicle System Dynamics* (2008) 46:6 (2008) 483-499.
- [8] P.T De Boer, D.P. Kroese, S. Mannor, R.Y. Rubinstein, A tutorial on the cross-entropy method. *Annals of Operations Research* 134 (2005) 19-67.
- [9] A. Belay, E.J. OBrien, D. Kroese, Truck fleet model for design and assessment of flexible pavements. *Journal of Sound & Vibration* 311 (2008) 1161-1174.

[10] D.J. Cole, D. Cebon, Truck tires, Suspension design and road damage, International Rubber Conference IRC'96, Manchester, UK, 1996.

[11] International Organization for Standardization ISO 8608, Mechanical vibration - road surface profiles - reporting of measured data, 1995.

[12] G.K. Chang, J.C. Dick, R.O. Rasmussen, *ProVAL Users guide*. The Transtec Group, Inc., 2006.

Figure 1 CE combinatorial optimisation algorithm process

Figure 2 Progression of CE algorithm for simple mathematical function with two maxima (—), showing samples (o) and the sampling distribution (----) for successive generations of the algorithm

Figure 3 Half-car roll vehicle model

Figure 4 Comparison of excitation inputs for half-car optimisation: Sine bump with $\lambda = 0.25$ m (+), sine bump for $\lambda = 0.5$ m (*), sine bump for $\lambda = 1.0$ m (o), sine bump for $\lambda = 1.5$ m (x), sine bump for $\lambda = 2.0$ m (\square), class A road profile (\diamond)

Figure 5 Observed (---) and difference between observed and best predicted (—) unsprung mass accelerations for vehicle model

Figure 6 Actual road profiles used for numerical validation (—) left wheel path (---) right wheel path (a) AR 1 (b) AR 2 (c) AR 3 (d) AR 4 (e) AR 5

- Figure 7 Estimated road profiles versus distance (—) left wheel path (---) right wheel path, and errors in estimation (o) left wheel path (Δ) right wheel path (a) AR 2 (b) AR 4
- Figure 8 Actual (—) and estimated (o) PSDs of road heights for left wheel of half-car (a) AR 1 (b) AR 2 (c) AR 3 (d) AR 4 (e) AR 5
- Figure 9 Moving average error in estimated profile from noise-free accelerations in left wheel path (—), from corrupted accelerations in left wheel path (—), from noise-free accelerations in right wheel path (---), from corrupted accelerations in right wheel path (---) for half-car algorithm. Left wheel path (solid) and right wheel path (dotted), noise-free (thin) and corrupted data (thick), (a) AR 2 (b) AR 4
- Figure 10 Actual (—), estimated (o) and estimated with noise (\bullet) PSDs for left wheel of half-car (a) AR 2 (b) AR 4
- Table 1 Vehicle model parameters
- Table 2 Details of road profiles used for numerical validation
- Table 3 Comparison of estimated and true IRI values for each profile
- Table 4 Comparison of estimated (from corrupted measurements) and true IRI values

Given target data set, γ_{obs} , and performance function $S(\chi) = \gamma$, determine the optimal parameters, χ^* , that yield $S(\chi^*) \approx \gamma_{\text{obs}}$

Initialisation

Define v_0 , a relatively arbitrary parameter set to initially describe the potential values of χ .

- Generate N_χ iid samples from $f(\chi | v)$

Optimisation

For $stages = 1, 2$

For $z = 1, 2, \dots, N_z$

- Evaluate $S(\chi_1), S(\chi_2), \dots$ etc. and compute the $(1-\rho)$ quantile
- Estimate v_z from the mean and variance of the corresponding $(1-\rho)$ quantile of samples of χ

End

After R iterations, reset v_z to a wider range of possible values to avoid local maxima. For this study, it is trebled.

End

Record the fittest solution $S(\hat{\chi}^*)$ and the corresponding parameter estimates \hat{v}^*

Fig. 1

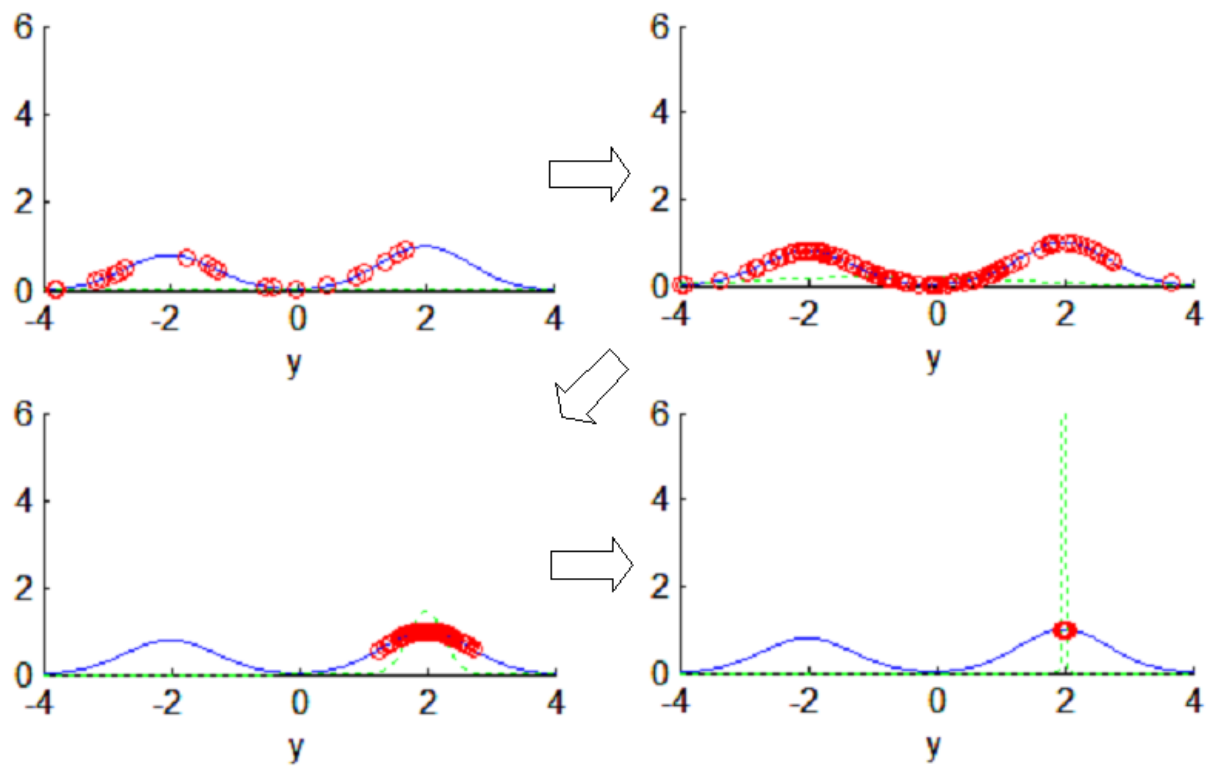


Fig. 2

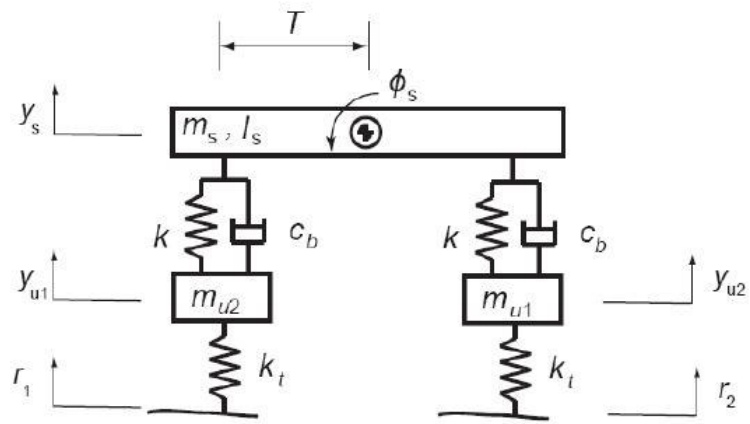


Fig. 3

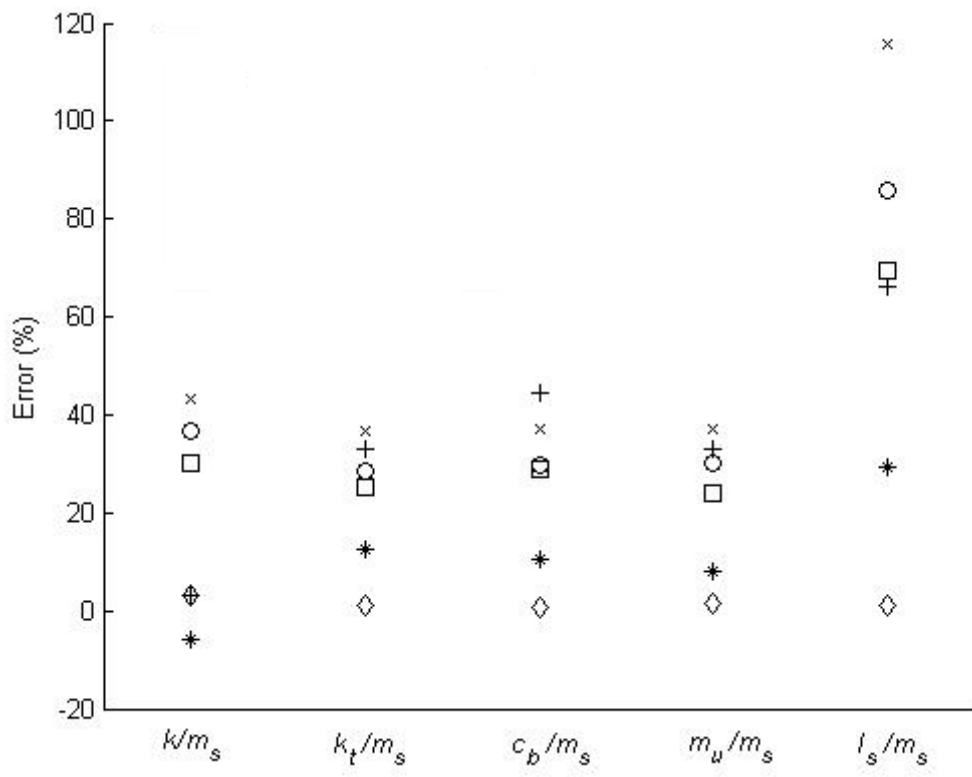


Fig. 4

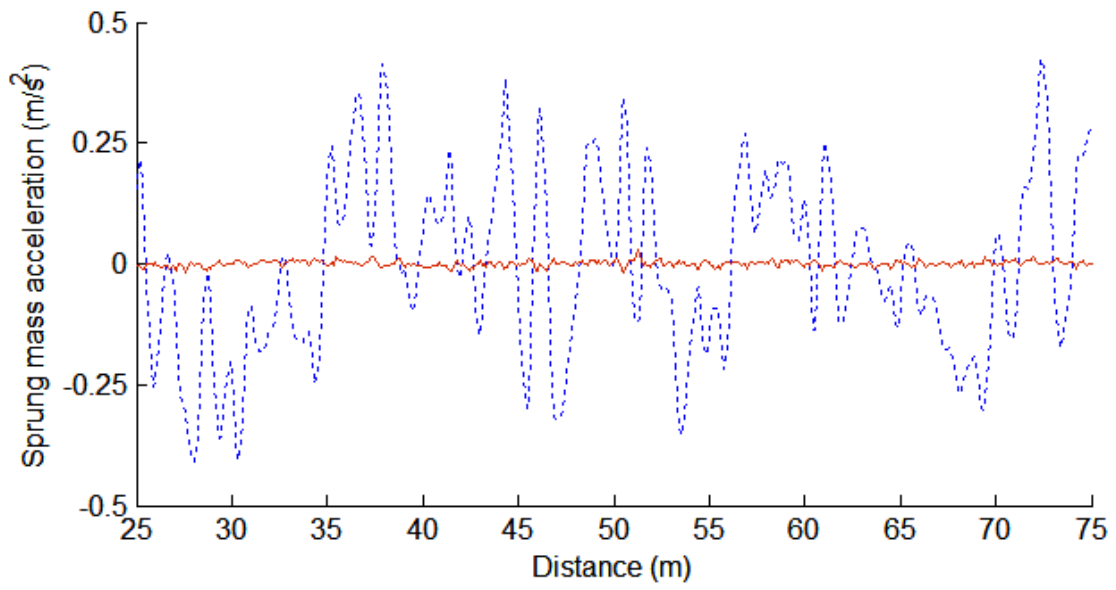


Fig. 5

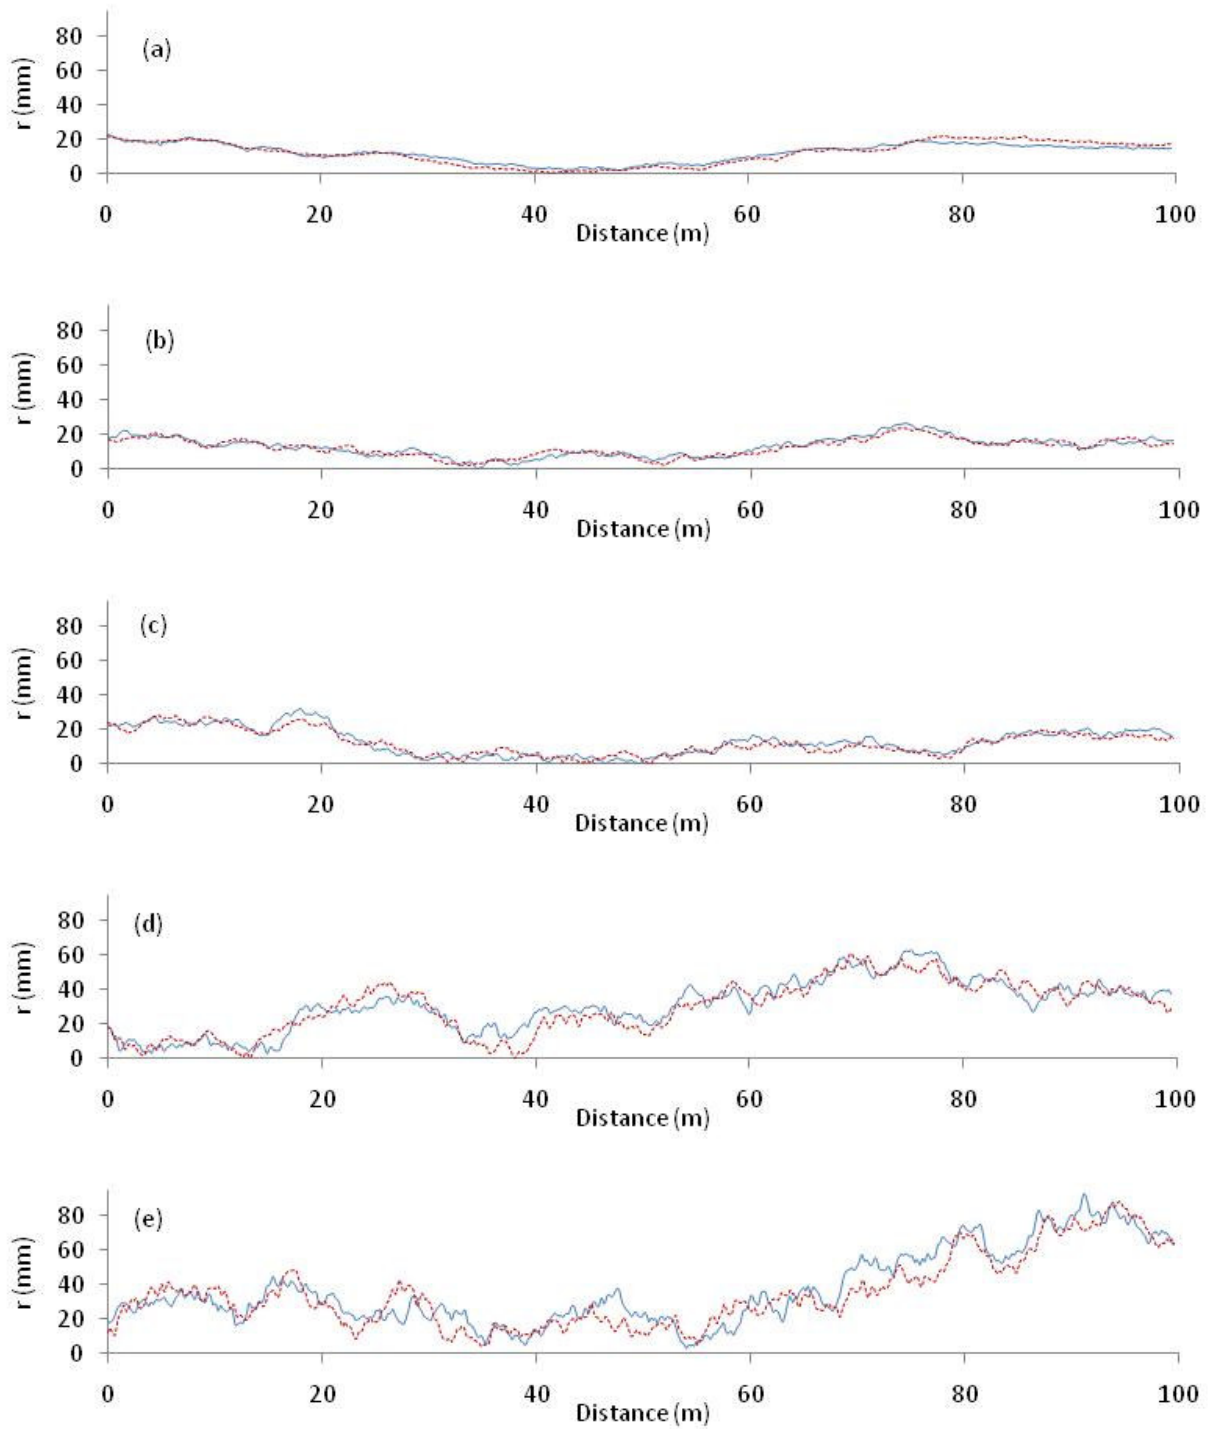


Fig. 6

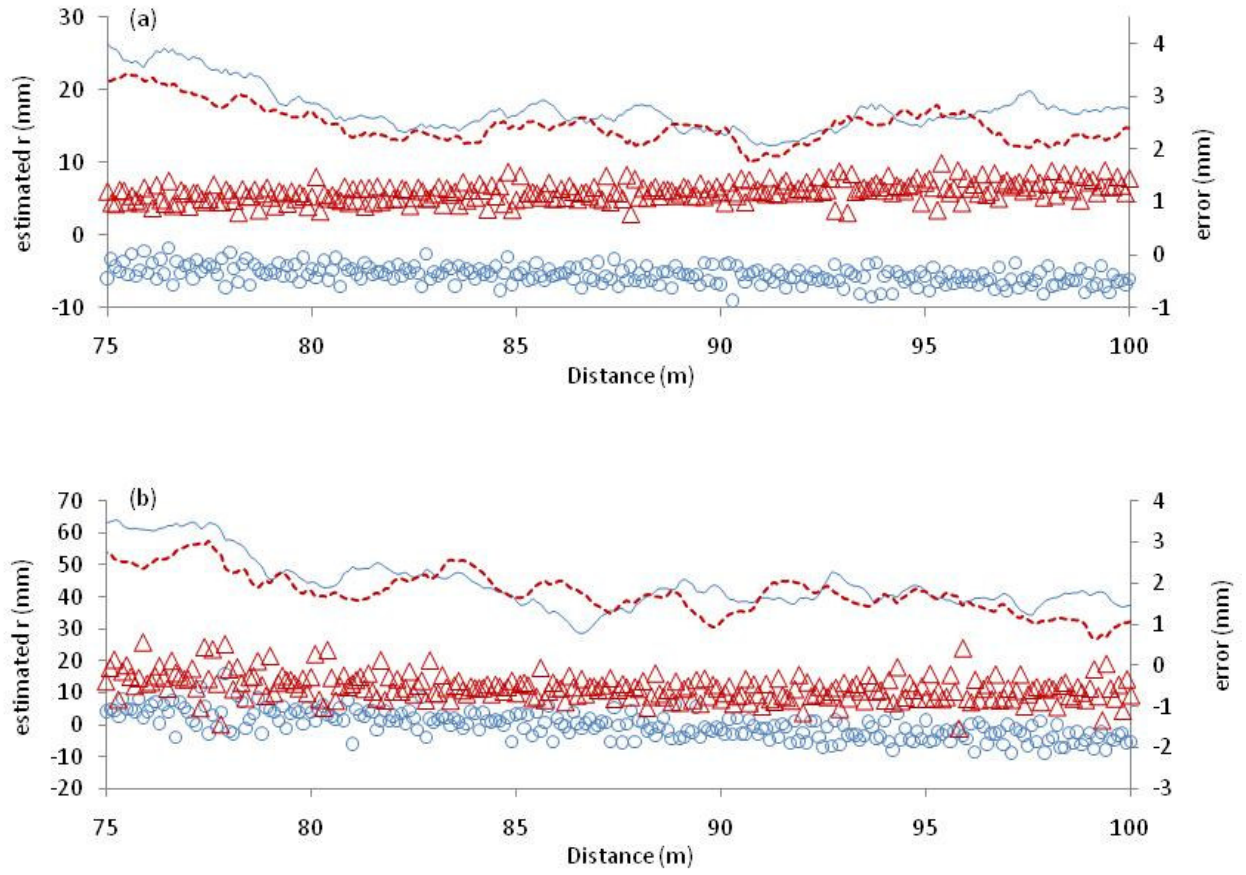


Fig. 7

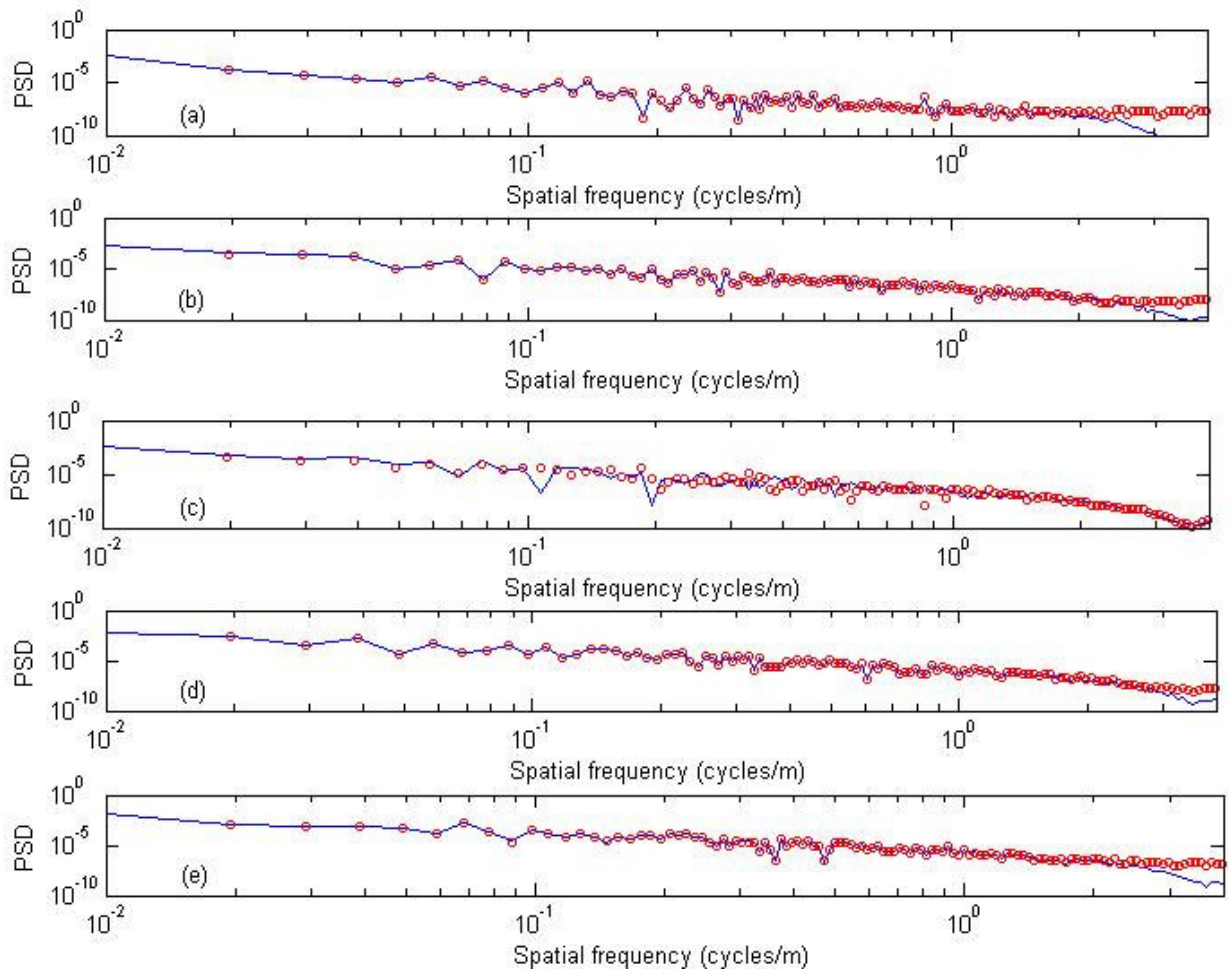


Fig. 8

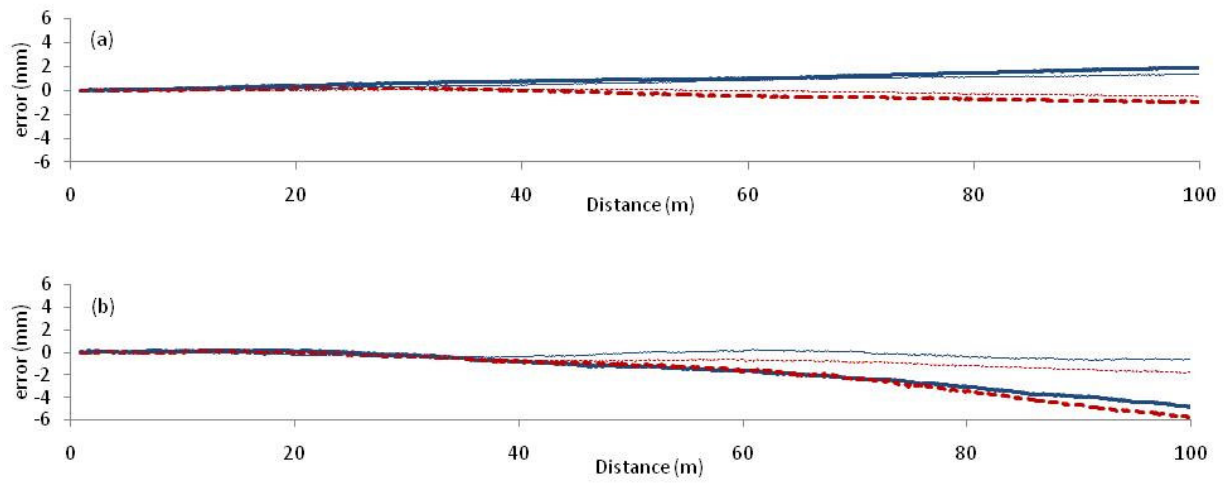


Fig. 9

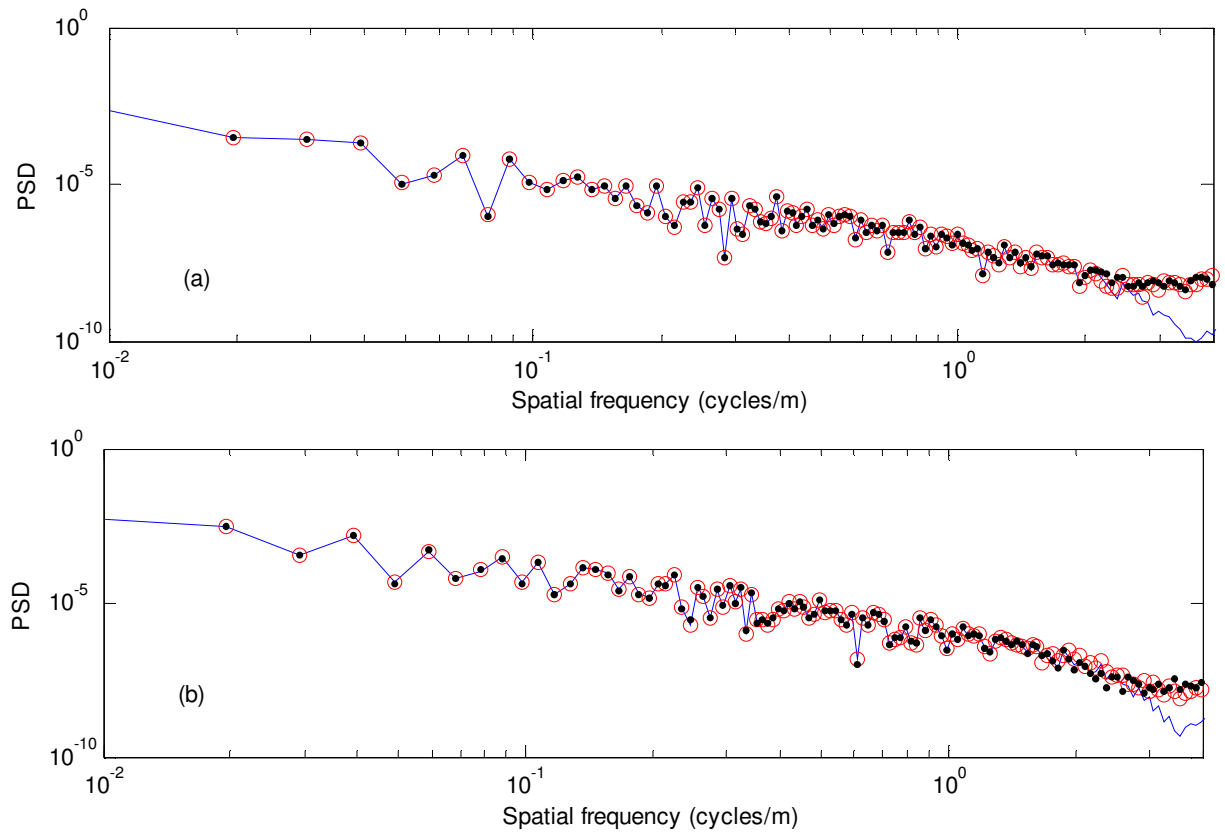


Fig. 10

Table 1

Description	Unit	Symbol	Value
Sprung mass	kg	m_s	500
Unsprung mass	kg	m_{u1}	25
		m_{u2}	25
Suspension stiffness	kN m ⁻¹	k	20
Tyre stiffness	kN m ⁻¹	k_t	150
Passive damping coefficient	kN s m ⁻¹	c_b	1.5
Roll moment of inertia	kg m ²	I_r	750
Track width (known parameter)	m	T	0.75

Table 2

	ISO Class	$G_d(n_0)$, $\text{m}^3 \text{ cycle}^{-1}$	IRI, m km^{-1} (left/right)		ρ_r
AR 1	Class A	8×10^{-6}	1.07	1.02	0.968
AR 2	Class A	16×10^{-6}	2.27	2.11	0.931
AR 3	Class B	64×10^{-6}	3.09	2.96	0.936
AR 4	Class B/Class C	128×10^{-6}	6.18	6.28	0.922
AR 5	Class C	256×10^{-6}	8.57	8.29	0.926

Table 3

<i>Left/right wheelpaths</i>				
Profile no.	IRI _{left} (est)	er (%)	IRI _{right} (est)	er (%)
AR 1	1.08	0.9%	1.03	1.0 %
AR 2	2.28	0.4 %	2.11	0 %
AR 3	3.10	0.3 %	2.97	0.7 %
AR 4	6.21	0.5 %	6.33	0.8 %
AR 5	8.52	0.6 %	8.33	0.5 %

Table 4

<i>Left/right wheelpaths</i>				
Profile no.	IRI _{left} (est)	er (%)	IRI _{right} (est)	er (%)
AR 2	2.28	0.4 %	2.11	0 %
AR 4	6.22	0.6 %	6.32	0.6 %