

Chapter 13

Making Structural Condition Diagnostics Robust to Environmental Variability

Harry Edwards, Kyle Neal, Jack Reilly, Kendra Van Buren, and François Hemez

Abstract Advances in sensor deployment and computational modeling have allowed significant strides to be made recently in the field of Structural Health Monitoring (SHM). One widely used SHM technique is to perform a vibration analysis where a model of the structure’s pristine (undamaged) condition is compared with vibration response data collected from the physical structure. Discrepancies between model predictions and monitoring data can be interpreted as structural damage. Unfortunately, multiple sources of uncertainty must also be considered in the analysis, including environmental variability and unknown values for model parameters. Not accounting for uncertainty in the analysis can lead to false-positives or false-negatives in the assessment of the structural condition. To manage the aforementioned uncertainty, we propose a robust-SHM methodology that combines three technologies. A time series algorithm is trained using “baseline” data to predict the vibration response, compare predictions to actual measurements collected on a potentially damaged structure, and calculate a user-defined damage indicator. The second technology handles the uncertainty present in the problem. An analysis of robustness is performed to propagate this uncertainty through the time series algorithm and obtain the corresponding bounds of variation of the damage indicator. The uncertainty description and robustness analysis are both inspired by the theory of info-gap decision-making. Lastly, an appropriate “size” of the uncertainty space is determined through physical experiments performed in laboratory conditions. Our hypothesis is that examining how the uncertainty space changes in time might lead to superior diagnostics of structural damage as compared to only monitoring the damage indicator. This methodology is applied to a portal frame structure to assess if the strategy holds promise for robust SHM.

Keywords Structural health monitoring • Uncertainty • Time series modeling • Robustness

13.1 Introduction

In the discipline of Structural Health Monitoring (SHM), the condition of a structure is often assessed using a combination of measurements and numerical models that attempt to infer the structural condition through changes in the vibration signature response. A structure inevitably degrades as it ages, through cycles of thermal loading, structural loading, and other conditions that affect structural health. At some point the structure may no longer be able to meet its performance requirements. This work contributes to the vast body of literature that develops diagnostics to detect such changes before they become safety or mission critical [1, 2].

H. Edwards

Atomic Weapons Establishment, Environment and Test Group, Reading, UK

e-mail: Harry.Edwards@awe.co.uk

K. Neal

Department of Civil and Environmental Engineering, Vanderbilt University, Box 1831-B, Nashville, TN 37235, USA

e-mail: kyle.d.neal@vanderbilt.edu

J. Reilly

Department of Civil and Environmental Engineering, Princeton University, Princeton, NJ 08544, USA

e-mail: jpr2@princeton.edu

K. Van Buren (✉)

Los Alamos National Laboratory, XCP-8, Mail Stop F644, Los Alamos, NM 87545, USA

e-mail: klvan@lanl.gov

F. Hemez

Los Alamos National Laboratory, XTD-IDA, Mail Stop T087, Los Alamos, NM 87545, USA

e-mail: hemez@lanl.gov

To detect the onset of structural changes that can adversely affect structural integrity, sensors are attached to the structure to measure its response from external excitation. Commonly, accelerometers and strain gauges are used to measure the structure's response. For in-service structures, ambient excitations such as cars on a bridge, wind over a building, or ground vibrations are typically used to elicit vibration response. In contrast, it is typical to perform deliberate and controlled excitation such as a modal hammer impact strike or modal shaker for laboratory tests. It is common practice to analyze the measurements through changes to the Frequency Response Function (FRF) data, which is appropriate for dynamics that remain mostly linear and stationary [3]. Another approach is to train a time series model on the pristine condition of the structure, and then use the model to predict the structure's response and assess if its condition remains unchanged. Shifts in natural frequencies observed from the FRF data, or deviations from predictions of the time series model, could indicate structural damage. It is emphasized that this process is only effective if measurements of the structure in its current state can be compared to measurements obtained from the pristine structure, which form a known "baseline." In the absence of baseline test data, mathematical modeling can be substituted to create theoretical "data" of the pristine condition.

In this study, a time series model from the family of Auto Regressive (AR) representations [4] is used to analyze the acceleration response data (measurements) from an aluminum frame structure tested in controlled laboratory settings. The AR model is trained using multiple sets of measured vibration response collected while the structure is in a "pristine," or undamaged, condition. The hypothesis is that the occurrence of damage manifests itself as a significant difference between what is measured on the (now damaged) structure and what the trained model predicts. The difference in structural state is diagnosed by the large prediction residuals that occur when the time series model is fit to damaged data that start to differ from the pristine data. A damage indicator is proposed that derives from statistics of the prediction residuals.

While the aforementioned strategy is relatively well accepted within the SHM community, one challenge, which remains for the most part unresolved, is the *management of uncertainty*. A change in vibration signature does not necessarily indicate that the structure is damaged. Instead, it might be due to environmental variability, such as a change in input excitation. To avoid the possibility of generating false-positives or false-negatives, it is highly desirable to "separate" the effects of environmental variability from those of structural damage.

Another unavoidable source of uncertainty is lack-of-knowledge that originates from the model used to predict vibration response. For example, when using the AR representation, the choice of the order of the model has been shown to affect the assessment of structural health [5]. Further, when training a model of arbitrary mathematical form, the model parameters might be non-unique whereby multiple sets of parameter values are able to replicate the training data with comparable fidelity. It is highly desirable to "isolate" the effects of model-form uncertainty from those of structural damage. We quantify the effects of environmental uncertainty by replicating the vibration tests as the input excitations are varied within reasonable bounds. Likewise, we quantify the effects of modeling uncertainty by analyzing, not a single best-fitted time series model, but a family of models that includes all representations that fit the measurements with a similar level of accuracy. The effects of these two main sources of uncertainty (environmental variability and model-form uncertainty) are quantified such that the detection of structural damage can be rendered *robust* to their unavoidable occurrence.

The description of environmental variability and modeling lack-of-knowledge is inspired by the theory of information-gap (or info-gap) for decision-making [6]. The magnitude, or "size," of the uncertainty, which the decision ("*is the structural state pristine or damaged?*") must be robust to, is controlled by a horizon-of-uncertainty parameter denoted " α ." A larger value of α increases the uncertainty considered in the analysis and, therefore, allows for greater potential deviation between reality and the numerical model used to predict the structural state.

This study aims to assess the structural state by examining the value of α needed to make the nominal AR time series model, which is developed to characterize the pristine state, predict a set of measurements obtained on a potentially damaged structure. The preliminary results indicate the success of this strategy. Damage states can be reliably diagnosed from observing the "growth" of the uncertainty space, that is, the extent to which α increases, while offering robustness to the environmental variability and modeling lack-of-knowledge.

13.2 Experimental Procedure

The test structure considered is a two-story portal frame made of aluminum bars and steel brackets with bolted connections as shown in Fig. 13.1a (left). The structure is clamped to the support table with four C-clamps. A modal shaker is bolted to one of the columns at 40.7 cm above the base plate. The structure is excited with a chirp sine wave whose frequency varies from 1000 Hz down to 2 Hz, and to which uniform white noise is added. It is verified that the applied chirp signal excites most of the lower resonant frequencies of the structure, which helps to capture a more descriptive response envelope. White noise added to the input shaker signal adds a variability of known magnitude to the excitation. It introduces one of the sources of environmental variability considered in the problem; the other sources are described later.




Fig. 13.1 Portal frame structure instrumented to perform the vibration tests. (a) *Left*: Test structure. (b) *Right*: Sensor placement

Table 13.1 Definition of structural states (pristine, variability, damage) considered

Case	Description	Variability?	Damage?	Replicates
0 ^(I)	First set of pristine (before cases I to V)	Yes	No	60
0 ^(II)	Second set of pristine (after cases I to V)	Yes	No	20
I	Table kicked during vibration testing	Not sure	Not sure	20
II	One C-clamp added to middle plate	No	Yes	20
III	Two C-clamps added to middle plate	No	Yes	20
IV	Loosening of top-right connection bolt	No	Yes	20
V	Thinner, lighter bottom-right column	No	Yes	20
VI	Removal of three base-plate bolts	No	Yes	20

To capture the response, four accelerometers are attached to the structure, as indicated by the red box in Fig. 13.1a (left). A detail is provided in Fig. 13.1b (right), where the green box shows the location of one of the accelerometers. Two sensors are located such that they can measure acceleration in the same direction as that of shaker input, while the other two are oriented to capture the out-of-plane response. Only the data collected at the location of the green box in Fig. 13.1b (right) are used for the analysis reported herein.

Vibration tests are performed over several days to accumulate sixty sets of pristine response data. Slightly different temperatures and ambient vibrations in the testing environment result from these combinations of days and times. This incorporates further environmental variability in the vibration response of the pristine condition. Our contention is that assessing damage should be made as immune as possible, or robust, to the presence of this environmental variability.

Two independent sets of vibration tests are executed to characterize the pristine condition. The first set is performed before we start to modify the portal frame to introduce different damage scenarios. The second set is performed after testing four of the damage cases. Assessing the statistical consistency of the response provides confidence that the “baseline” (pristine) state of the structure has not significantly shifted due to the various tests performed.

Damage is introduced in the structure in a number of different ways, as summarized in Table 13.1. Firstly, mass is added in the form of one and then two C-clamps attached to the middle plate. These cases are labeled II and III in Table 13.1. A single C-clamp weighs about 1 % of the total mass of the portal frame. The mass loading that results is believed to induce a linear change in the vibration response of the structure. Secondly, the bolt located in the top-right connection is loosened (Case-IV). This case potentially generates nonlinear vibrations from the impact between the vertical column, horizontal plate, and

attachment bracket of the portal frame. An advantage of the time series modeling pursued in this work is that it can be used to detect linear and nonlinear changes to the response alike. Thirdly, Case-V replaces the bottom-right vertical column with a thinner (skinnier) and lighter column and Case-VI removes three of the four bolts that connect the vertical column to the horizontal base plate. While its response remains linear, Case-V defines a significant structural change. Case-VI introduces a potential non-linearity in the vibration response because of the higher likelihood for column-to-base-plate impact.

Table 13.1 also includes Case-I that keeps the frame structure in its pristine condition, identical to cases $0^{(I)}$ and $0^{(II)}$, while a member of the team “kicks” the support table during the execution of vibration tests. This unorthodox procedure is meant to define an extreme case of environmental variability where the excitation provided to the portal frame is not even stationary. The table lists the numbers of replicates performed for each case. It can be observed that this experimental campaign results in a significant number of vibration tests ($60 + 7 \times 20 = 200$ tests).

The third and fourth columns of Table 13.1 attempt to categorize the perturbations introduced to the frame structure as either “variability” or “damage.” The mass loading cases (II and III) are meant to represent, for example, a heavy piece of equipment that would be installed inside a building. Loosening a connecting bolt defines a clear damage scenario. It is unclear how Case-I, and its extreme form of exciting the structure by kicking the support table, should be categorized. Our main objective is to demonstrate that the damage cases (II, III, IV, V, and VI) can be separated from the pristine cases ($0^{(I)}$ and $0^{(II)}$) in the presence of environmental variability represented by replicate testing of the pristine structure across multiple days. Another success would be to demonstrate that Case-I is not confused for structural damage.

13.3 Diagnosis of Structural Damage in the Presence of Environmental Variability

The procedure being implemented to diagnose the occurrence of structural damage from vibration measurements is to train a time series model using “baseline” data that represent the pristine structural state. Baseline model predictions are compared to actual measurements collected on a potentially damaged structure, and a damage indicator is calculated. This section summarizes the procedure and explains how the experimental variability is handled.

We start by considering three time series models for analysis of the raw acceleration signals: an Auto Regressive (AR) model, an Auto Regressive Moving Average (ARMA) model, and an Auto Regressive Exogenous (ARX) model [7]. Equation 13.1 defines the generic form of these models, which represents the measurements using a linear combination of input and output responses:

$$\hat{y}_k = \sum_{1 \leq j \leq N} \beta_j \cdot y_{k-j} + \sum_{1 \leq j \leq M} \gamma_j \cdot f_{k-j} + e_k \quad (13.1)$$

where f_k denotes the input excitation (force applied by the shaker) at time sample t_k , y_k denotes the output response (acceleration measured by the accelerometer), and e_k is a statistical white-noise term that accounts for model-fitting errors. The symbol \hat{y}_k is the prediction of the model, which can be compared to the actual measurement, y_k , to define the prediction residual, ε_k :

$$\varepsilon_k = y_k - \hat{y}_k \quad (13.2)$$

In Eq. (13.1), the optimal values of orders (N; M) of AR and MA summations are determined by repeating the training procedure, as discussed in the following. Once the AR and MA model orders are decided, the regression coefficients (β_j ; γ_j) are fitted to the “baseline” data.

In contrast to Eq. (13.1), an AR model is restricted to the first summation that defines a linear combination of N past observations (y_{k-1} ; y_{k-2} ; ... y_{k-N}). An ARMA model is comprised of the AR term to which a contribution that represents an exogenous input is added. In the case of an unknown excitation signal, this additional term can be defined as the residual (13.2) obtained from AR model predictions. These residuals ε_k are then substituted to the input f_k in Eq. (13.1).

An ARMA model is initially selected to model the vibration response of the frame structure. A study is undertaken to determine the optimal orders (N; M) of summations in Eq. (13.1). Both orders are varied to analyze all (N; M) combinations ranging from $1 \leq N \leq 25$ for the AR term and $1 \leq M \leq 30$ for the MA term. The mean absolute value of prediction residuals (13.2) is used as a metric for goodness-of-fit between actual measurements (y_k) and predictions (\hat{y}_k) of the model.

Figure 13.2 indicates how the mean statistics of residual error vary as a function of orders (N; M). These results are obtained for one of the replicate tests of the pristine structure and it is verified that similar trends are observed with the other replicates. A lower residual mean value describes a model that predicts the measured acceleration response more accurately.



Fig. 13.2 Mean values of ARMA residual errors for varying model orders (N ; M)

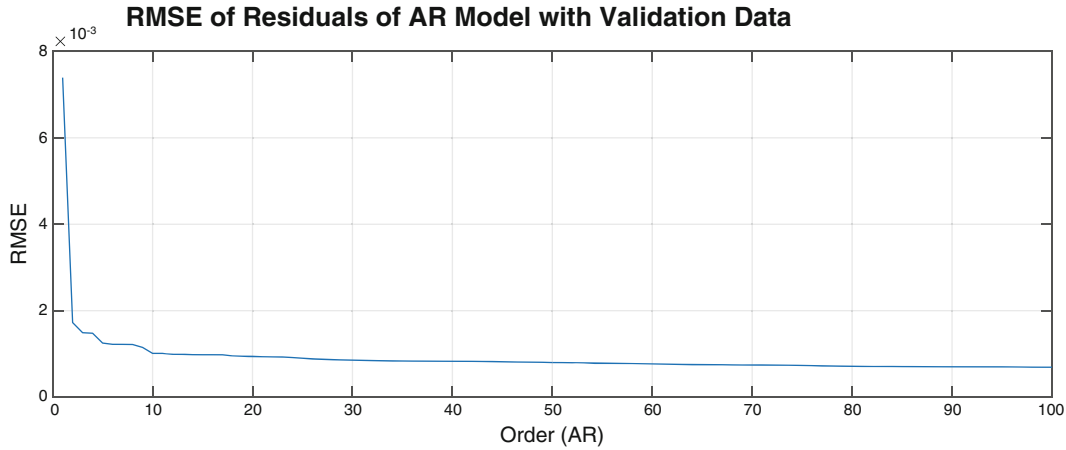


Fig. 13.3 Mean values of AR residual errors for varying model orders, $1 \leq N \leq 100$

From Fig. 13.2, it is clear that, first, an AR order of $N = 10$ suffices to “stabilize” the residual error and, second, the MA term does not seem to have a significant impact on the goodness-of-fit.

In light of the results illustrated in Fig. 13.2, it is decided to conduct our investigation with an AR model. Figure 13.3 shows the Root Mean Square (RMS) value of AR prediction errors (13.2) for model orders that range in $1 \leq N \leq 100$. The figure indicates that a model order of $N = 10$ is sufficient to reach an acceptable level of prediction accuracy while avoiding over-fitting the measurements. For completeness, it is mentioned that ARX models are investigated in a similar manner, with no improvement in accuracy over that of AR models. Results discussed in the remainder of the manuscript are obtained with AR models of order $N = 10$.

After having selected the modeling approach (AR) and model order ($N = 10$), the ten unknown regression coefficients ($\beta_1; \dots \beta_{10}$) of Eq. (13.1) are best-fitted to acceleration measurements for each replication of Case-0^(l), see Table 13.1. The “baseline” AR model of the pristine condition is defined by averaging the regression coefficients across replicate tests:

$$\langle \beta_j \rangle = \frac{1}{R} \cdot \sum_{1 \leq r \leq R} \beta_j^{(r)} \quad (13.3)$$

where the superscript $()^{(r)}$ identifies one of the replicate tests, that is, $1 \leq r \leq R$, and brackets $\langle \rangle$ indicate an averaged value. Figure 13.4 illustrates this procedure, which yields the averaged coefficients, $\langle \beta_1 \rangle \dots \langle \beta_{10} \rangle$, of the “baseline” AR model.




Fig. 13.4 Definition of the “baseline” AR model from tests of the pristine structure




Fig. 13.5 Damage indicator defined from the vibration test of a different structural state

Once the “baseline” AR model is available to characterize the response of the pristine condition, it is analyzed to predict the structural state given a new set of experimental vibration data. The predictions are compared to actual measurements to calculate the residual errors (13.2). The mean value and standard deviation value of prediction residuals are estimated and used to define our Damage Indicator (DI) as a combination of the two statistics:

$$DI = \left(\frac{\mu}{\mu_0} \right) + \left(\frac{\sigma}{\sigma_0} \right) \quad (13.4)$$

where μ and σ denote the mean and standard deviation values, respectively. In (13.4), the symbols μ_0 and σ_0 are “reference” values used to scale the two contributions to DI. Scaling is needed because the mean statistics are observed to be several orders of magnitude smaller than the standard deviation values. Figure 13.5 illustrates the procedure to calculate DI.

The sixty replicate tests defined for Case-0^(l) are divided into training and validation sets. Fifty tests are used for training the “baseline” AR model, that is, $R = 50$ in Fig. 13.4. The remaining ten tests are utilized to assess how the model performs to predict data that have not been used for training, which is the validation step. The “reference” statistics (μ_0 ; σ_0) of Eq. (13.4) are estimated from the residual errors of these ten validation tests. Since replicates of the Case-0^(l) pristine condition are slightly different because of environmental variability (see Sect. 13.2), this procedure is deemed appropriate to capture its effect on predictions of the “baseline” AR model.

The “baseline” model is defined from multiple replicates of the pristine state. Thus, generating the DI as shown in Fig. 13.5 with a new time series measurement collected from the structure in a pristine condition should result in a low magnitude of residual error. The corresponding value of the DI should be close to 2, which would represent similar statistics as those obtained when analyzing the measurements used for training. Processing a time series measurement that represents a different structural state, on the other hand, should generate prediction residuals whose magnitude exceeds that produced by environmental variability alone. The value of the DI metric should increase. This is how we expect to separate the effect of structural damage from that of environmental variability, irrespective of the vibration dynamics (linear or nonlinear).

After training the “baseline” AR model to represent the pristine condition, the process of Fig. 13.5 is applied to each one of the twenty replicate tests of Case-I, Case-II, Case-III, and Case-IV, see Table 13.1. Within any one of these configurations, the excitation signal applied to the structure differs slightly from test to test due, as mentioned earlier, to the addition of random white noise. It means that the vibration response of the portal frame varies slightly as well for each test. This variability propagates to the prediction residuals, their statistical moments, and the DI metrics.

Figure 13.6 represents the distributions of replicate DI values obtained for the four perturbed states as ranges that extend from the minimum to maximum values. Each color (red, magenta, green, and black) corresponds to one of cases I to IV. The blue interval indicates the range of DI values that results from replicates of the Case-0^(I) pristine condition. Black crosses denote the average values of the test replicates.

For each configuration tested, the ranges of DI values capture the effects of the environmental variability. The question is whether the structural perturbations of Cases I to IV produce shifts in DI that can be “separated” from the aforementioned ranges. This is clearly observed for Case-II (magenta interval) and Case-III (green interval), which perturb the nominal configuration of the frame structure through mass loading. The conclusion is similar for Case-IV (black range) where one of the column-to-plate connecting bolts is loosened. The ranges obtained for these three configurations are completely outside the DI bounds obtained for the pristine structure. These cases, therefore, can be confidently classified as damaged. It is interesting to notice that cases II and III register such large DI shifts, even though the mass added is modest (1–3 %).

The red interval, representing the artificial “quake” of Case-I where the support table is kicked during vibration testing, cannot be unambiguously identified as either damaged or undamaged. This is because the interval partially overlaps with the DI bounds of the pristine structure, as indicated in Fig. 13.6. This perturbation represents a case of extreme variability but

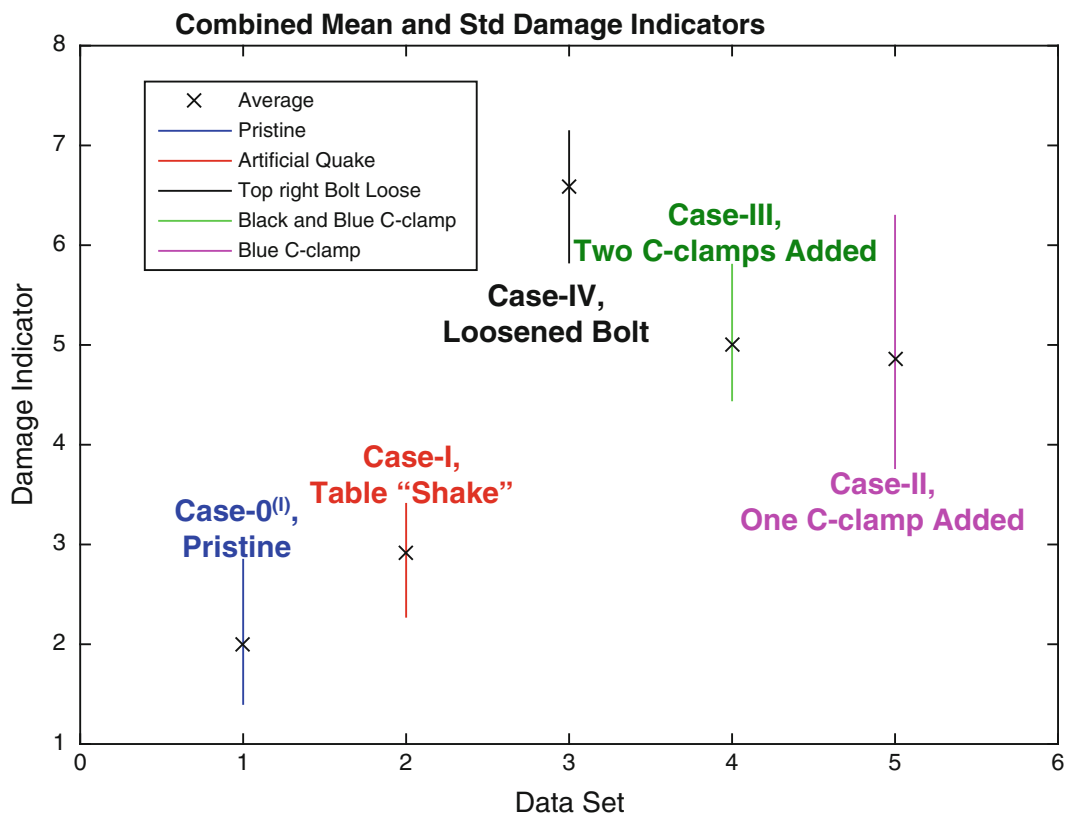


Fig. 13.6 Damage indicators of Case-0^(I) (pristine) and cases I to IV (perturbed)

not damage. Because the training data do not encompass such a significant variability, the DI bounds of the pristine state (blue interval) do not fully overlap with DI values obtained for Case-I.

The artificial “quake” of Case-I stands out as a condition that our procedure cannot classify well. The preliminary results obtained nevertheless demonstrate the strong potential that this method offers to classify with accuracy the structural condition despite environmental variability.

13.4 Rendering the Damage Assessment Robust to Modeling Lack-of-Knowledge

In this section, we discuss how to render the assessment of structural condition robust to the inevitable modeling lack-of-knowledge. This modeling uncertainty, here, is represented by the unknown regression coefficients, β_k , of the AR model. Even though it is not discussed here, other sources of modeling uncertainty could easily be included in the proposed methodology. The AR model order, for example, could be considered to be a modeling uncertainty as well as the functional form of the time series representation (ARMA, support vector machine, other?).

The objective of the proposed methodology is to guarantee that the assessment of structural condition (pristine, damaged?) is as robust as possible to the fact that parameters of the time series model are not known precisely. We would not want to find ourselves in a situation where, for example, the assessment of structural condition changes from damaged to undamaged, or vice-versa, simply because the values of model parameters shift slightly due to uncontrolled factors. Establishing robustness provides confidence that the decision reached (“*the structure is damaged*”) is correct despite the fact that model parameters vary from their nominal values.

To establish the robustness of condition monitoring, each regression coefficient of the “baseline” AR model is varied between a lower bound, β_k^{Lower} , and an upper bound, β_k^{Upper} , as:

$$\langle \beta_k \rangle + \alpha \cdot (\beta_k^{\text{Lower}} - \langle \beta_k \rangle) \leq \beta_k \leq \langle \beta_k \rangle + \alpha \cdot (\beta_k^{\text{Upper}} - \langle \beta_k \rangle) \quad (13.5)$$

where α denotes the dimensionless, “horizon-of-uncertainty” parameter greater or equal to zero. The lower and upper bounds are obtained from the population of model regression coefficients ($\beta_k^{(1)}$; $\beta_k^{(2)}$; ... $\beta_k^{(R)}$) calculated in Fig. 13.4 during the training step of the pristine condition.

Equation 13.5 varies each regression coefficient, β_k , between the lower and upper bounds, that is, $\beta_k^{\text{Lower}} \leq \beta_k \leq \beta_k^{\text{Upper}}$, depending on the value of parameter α . With $\alpha = 0$, there is no variation of the regression coefficient that defaults back to its “baseline” value, that is, $\beta_k = \langle \beta_k \rangle$ as before. With $\alpha = 1$, the regression coefficient can take any value between the lower and upper bounds.

The procedure starts by setting the horizon-of-uncertainty parameter α , then, the assessment of structural condition of Fig. 13.5 is repeated multiple times by varying the values of regression coefficients within the lower and upper bounds of Eq. (13.5). Each calculation of the DI metric is based on predictions of an AR model with coefficients (β_1 ; ... β_{10}) that deviate from the mean-value coefficients ($\langle \beta_1 \rangle$; ... $\langle \beta_{10} \rangle$) of the “baseline” AR model. The magnitude of the deviation is controlled by α . The parameter α , therefore, defines the modeling uncertainty embodied, here, by unknown regression coefficients. The goal of the robustness analysis is to demonstrate that the assessment of structural condition can withstand as much modeling uncertainty as possible, that is, an as-high-as-possible value of α .

When performing the robustness analysis, however, the ten model parameters cannot be varied independently from each other. This is because they are statistically correlated. Figure 13.7 depicts the correlation structure for pairs of regression coefficients (β_p ; β_q), $p \neq q$. The figure indicates a strong correlation for all pairs of model parameters except combinations (β_7 ; β_8) and (β_8 ; β_9).

Figure 13.8 is a notional illustration of what happens when the correlation structure is ignored as the regression coefficients (β_1 ; ... β_{10}) are generated. The green-dot points remain “true” to the nominal correlation suggested by the ellipsoid shape (blue dashed line). While sampling the model parameters, we wish to avoid generating red-cross points that fall outside the ellipsoid. Independently varying each model parameter according to (13.5) would generate samples that do not accurately map the modeling uncertainty space of the “baseline” AR representation.

To maintain the correlation between regression coefficients of Fig. 13.7, a Principal Component Analysis (PCA) is implemented to transform the coefficients to a generalized coordinate space according to the decomposition:

$$\beta - \langle \beta \rangle = U \cdot \Sigma \cdot V^T \quad (13.6)$$

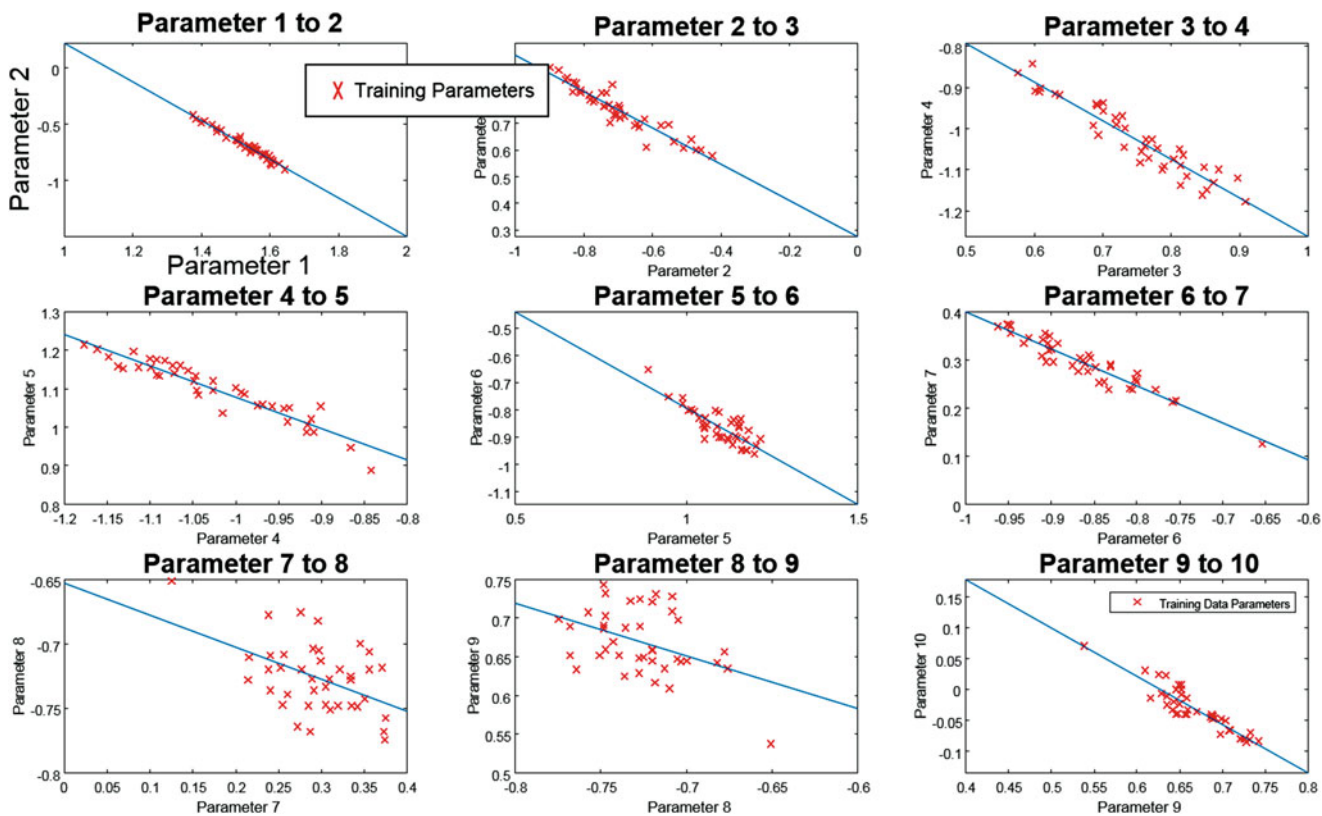
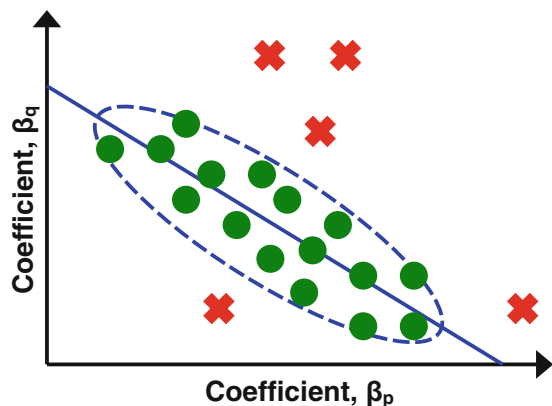


Fig. 13.7 Pairs $(\beta_p; \beta_q)$, $p \neq q$, of AR regression coefficients obtained for Case-0⁽¹⁾

Fig. 13.8 Coefficients sampled inside (green dot) or outside (red cross) the correlation



where the symbol β is the R-by-N matrix of regression coefficients, $\langle \beta \rangle$ is a matrix of the same size that removes the mean values $(\langle \beta_1 \rangle; \dots \langle \beta_{10} \rangle)$, and the triplet of matrices $(U; \Sigma; V)$ in the right-hand side is the PCA decomposition [8]. Equation 13.6 converts the correlated parameters $(\beta_1; \dots \beta_{10})$ of the AR model to zero-mean, uncorrelated generalized coordinates $(V_1; \dots V_{10})$.

After calculating the triplet of matrices $(U; \Sigma; V)$ of the decomposition (13.6), it is used “in reverse” to generate correlated samples of regression coefficients $(\beta_1; \dots \beta_{10})$, as illustrated notionally in Fig. 13.9. Because the generalized coordinates $(V_1; \dots V_{10})$ are uncorrelated, a Latin Hypercube Sample (LHS) is implemented to sample their values according to a uniform distribution [9]. The LHS method offers an efficient trade-off between randomness and spreading the samples throughout the 10-dimensional hypercube $([-1; +1])^{10}$ in which the generalized coordinates are defined. Equation 13.6 is used to transform the uniform samples to the correlated space whose principal directions are defined by the left singular vector $(U_1; \dots U_{10})$ of the PCA decomposition. This entire procedure is implemented in Matlab[®] and only takes a fraction of second to complete when implemented on a single-core desktop processor.




Fig. 13.9 Transform implemented to sample the regression coefficients ($\beta_1; \dots \beta_{10}$)

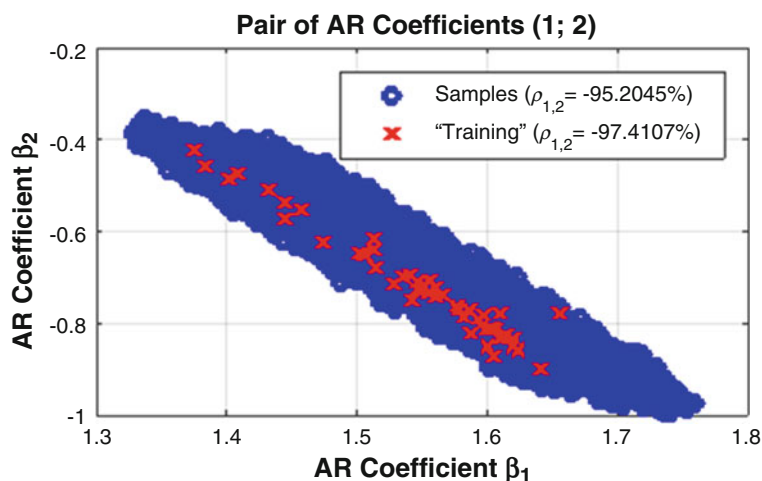


Fig. 13.10 Nominal and new correlation structure for the pair ($\beta_1; \beta_2$) of AR coefficients

A Monte Carlo simulation is analyzed using 10,000 sets of regression coefficients ($\beta_1; \dots \beta_{10}$) obtained using the combination of PCA transform and LHS method previously described. The β -values are consistent with the lower and upper bounds of Eq. (13.5) where $\alpha = 1$. A value of the DI is obtained for each Monte Carlo sample, following the flowchart of Fig. 13.5, and with the only difference that samples ($\beta_1; \dots \beta_{10}$) are substituted to the mean values ($\langle \beta_1 \rangle; \dots \langle \beta_{10} \rangle$).

Figure 13.10 depicts the 10,000 samples obtained for the pair of model parameters ($\beta_1; \beta_2$), which features the strongest correlation of all pairs ($\beta_p; \beta_q$), $p \neq q$. Red crosses represent the training data of the Case-0^(l) pristine condition and blue dots are the 10,000 new samples created from Eq. (13.5) at the level of modeling uncertainty of $\alpha = 1$. The figure indicates that the overall correlation is preserved by our sampling algorithm, as suggested by a coefficient of correlation between β_1 and β_2 equal to -97.4% for the training points and -95.2% for the new samples.

Figure 13.10 also suggests that the level of modeling uncertainty used, $\alpha = 1$, generates variations of regression coefficients ($\beta_1; \dots \beta_{10}$) that exceed those observed when replicating the vibration tests for the Case-0^(l) pristine configuration. The perturbations of model parameters generated in this manner account for variability not directly observed during vibration testing and that could manifest itself for that specific structural state at another time. In addition, these perturbations are appropriate to explore the modeling lack-of-knowledge, that is, the fact that adopting a different time series modeling approach or changing the model order(s) could result in a larger-than-currently-observed variation of model parameters. The robustness question is: *“do these variations of regression coefficients change the assessment of structural condition?”*

The above question is answered by running a Monte Carlo simulation of 10,000 LHS-generated AR models through each case defined in Table 13.1, including the training, support table “quake,” and damage cases. Figure 13.11 displays the ranges of DI values obtained from the analysis of each pristine and damage configuration (20 sets of data each) run through these 10,000 AR models. The difference with results shown in Fig. 13.6 is that the uncertainty considered here is much “broader” than the environmental variability investigated in Sect. 13.3.

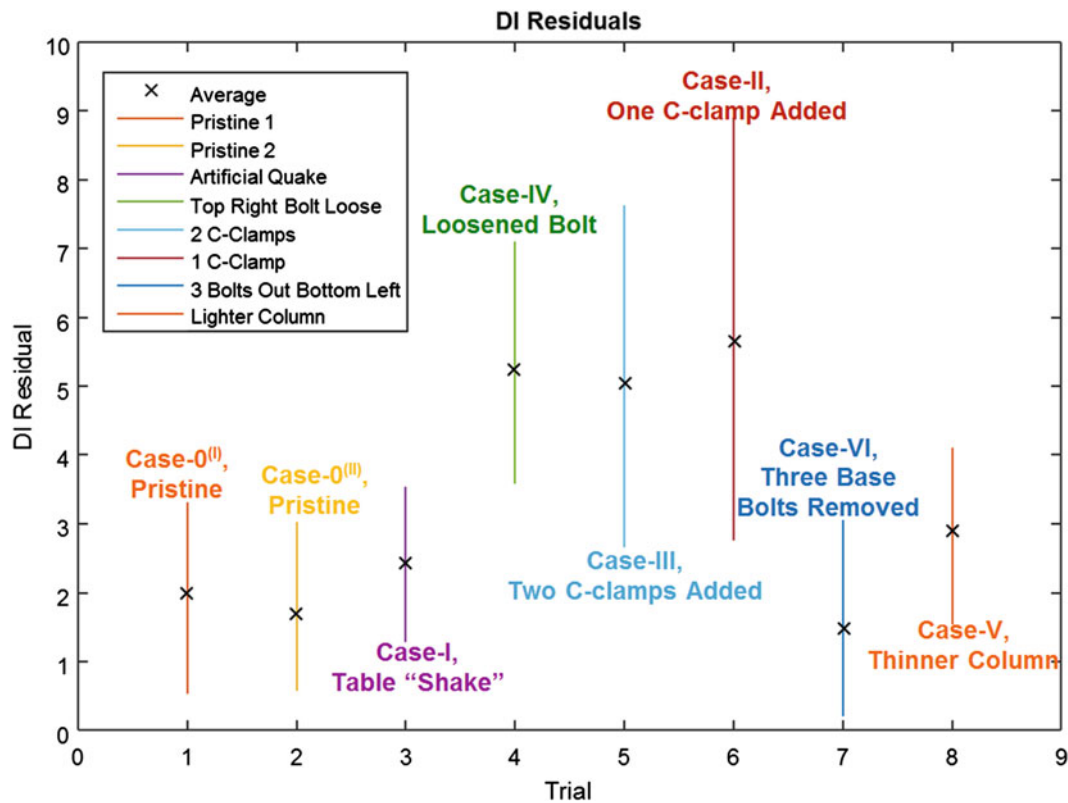


Fig. 13.11 Ranges of damage indicators that account for modeling uncertainty (at $\alpha = 1$)

It can be observed in Fig. 13.11, firstly, that the DI ranges for the two independent tests of the pristine condition, Case-0^(I) (orange interval) and Case-0^(II) (yellow interval), are remarkably consistent. This gives confidence that our procedure is not adversely affected by small changes to the “baseline” state that might have resulted from modifying the frame structure multiple times and testing it over-and-over. The second observation is that Case-I (purple interval), which represents the artificial “quake” where the support table is kicked during vibration testing, can no longer be separated from the pristine state. This result is welcome since this configuration, while it represents an extreme form of environmental variability, is not, strictly speaking, damaged.

The significance of this observation is emphasized. The consistency of DI bounds for the three sets, Case-0^(I) (orange), Case-0^(II) (yellow), and Case-I (purple), provides a “calibration” of the horizon-of-uncertainty parameter, α , needed to capture all pristine states of the frame structure. The submersion of the Case-I of extreme variability into the pristine bounds is predicated on a rather large number of trials (20 replicates by 10,000 samples = 200,000 trials), which supports the rigorousness of our exploration of the uncertainty space. The DI bounds predicted for these three states overlap, for the most part, when the modeling lack-of-knowledge represented by (13.5) is explored up to $\alpha = 1$.

Said differently, $\alpha = 1$ is the *smallest size* of the uncertainty space that produces consistent DI bounds for the pristine state. It means that the pristine condition of the frame structure can be *robustly diagnosed*, up to $\alpha = 1$, irrespective of the uncertainty of the vibration testing protocol (nominal testing vs. support table “shake”) and environmental variability (test-to-test replication).

The aforementioned adjustment of horizon-of-uncertainty, α , using the different pristine states of the portal frame, leads to a third observation in Fig. 13.11. When analyzed at the level of $\alpha = 1$, the added-mass and loosened bolt configurations are clearly identifiable. Their DI bounds (Case-II, magenta interval, Case-III, light blue interval, and Case-IV, green interval) mostly do not overlap with those of the first three pristine cases, even though the lack-of-knowledge of regression coefficients in Eq. (13.5) is propagated through the analysis.

Repeating the procedure with a larger horizon-of-uncertainty automatically produces larger DI bounds because Eq. (13.5) defines nested intervals of uncertainty as α increases. One could conceive, therefore, increasing the value of α until the DI




Fig. 13.12 Comparison of empirical histograms of DI values (at $\alpha = 1$)

bounds of the pristine and damaged configurations start to overlap. We estimate that reaching the point where the DI intervals completely overlap requires a horizon-of-uncertainty equal to $\alpha = 3$, approximately.¹

Because a modeling uncertainty space of size $\alpha \approx 3$ is unambiguously larger than the space of size $\alpha = 1$ inferred from “calibration” to the pristine configurations, it can be concluded with confidence that Case-II, Case-III, and Case-IV are damaged states. These results illustrate the potential offered by our strategy for robust SHM, whereby the “growth” of the uncertainty space (parameter α) is quantified as opposed to simply monitoring the value of a damage indicator.

Lastly, Fig. 13.11 shows that the procedure does not successfully “isolate” the configurations of Case-V (orange interval, thinner column) and Case-VI (dark blue interval, removal of three base bolts) from the pristine condition. These two damage cases represent slightly less severe alterations of the vibration dynamics than those of Case-II/Case-III (mass loading) and Case-IV (loosened bolt). We hypothesize that it is the reason why they remain undetected. In addition, these scenarios introduce structural perturbations at the bottom-left side of the portal frame, while the measurements analyzed are collected from a single sensor located at the top-right. Even though this conjecture has not been verified, sensor data from the other accelerometers, or the addition of more sensors, could have possibly provided a better classification.

For completeness, it is noted that our sampling-based exploration does not just produce the lower and upper bounds of the DI metric at any value of α . An entire population of DI values is available for each damage state, from which statistics can be estimated. Figure 13.12 illustrates the empirical distributions of DI metrics for Case-0^(I) (dark blue line, pristine), Case-III (magenta line, two C-clamps added), Case-IV (yellow line, loosened bolt), and Case-V (green line, thinner column).² This information is richer than the ranges shown in Figs. 13.6 and 13.11, which could lead to the development of probabilistic-based damage metrics. For example, the probability of falsely classifying a damaged configuration as pristine is defined by the overlap area of the damaged and pristine data. Figure 13.12 shows that these false-positive probabilities are 1.25 % for Case-III (mass loading), 0 % for Case-IV (loosened bolt), and 20 % for Case-V (thinner column).

These statistics do not alter the picture offered in Fig. 13.11 for the mass loading and loosened bolt cases, since the clear separation of their intervals from those of the pristine configurations already classifies them as damaged. The finer

¹In the interest of full disclosure, it is mentioned that $\alpha \approx 3$ is an estimate based on the assumption that the damage indicator grows linearly with α . While this rule-of-thumb is believed to be conservative, it has not been verified by performing a numerical optimization to search for the minimum and maximum values of DI over the uncertainty space defined in Eq. (13.5). (See References [10, 11] for further details.)

²The reader is warned that the color schemes shown in Figs. 13.6, 13.11, and 13.12 are not consistent.

resolution brought by estimating probabilities, however, does alter the conclusion for Case-V (thinner column). While it is indistinguishable from the pristine condition based on comparing DI bounds in Fig. 13.11, the 80 % probability of damage estimated from Fig. 13.12 strongly suggests that Case-V is different from pristine.

13.5 Conclusion

A commonly encountered strategy to deploy Structural Health Monitoring (SHM) solutions is to develop a mathematical or numerical model that represents a given condition, such as the pristine (undamaged) state, and compare predictions to actual data collected on the structure. Discrepancies between the two sets are interpreted as a change in condition and, potentially, structural damage. This strategy generally leaves unaddressed the difficult question of how to handle the sources of uncertainty that can include environmental variability, unknown functional forms of the mathematical models, and non-unique model parameter values.

To meet this challenge, we propose a framework where the assessment of structural condition is made *robust* to the aforementioned sources of uncertainty. Instead of monitoring the value of a damage index, we suggest to monitor the “growth” of the uncertainty space that describes the environmental variability and modeling lack-of-knowledge. An increasing level of uncertainty indicates that the structural condition deviates more-and-more from the nominal, “baseline” state and, therefore, is suspected to be evolving towards damage.

This novel concept is illustrated by testing a portal frame structure in different configurations that include pristine, mass-loaded, and damaged states. We demonstrate that the appropriate “size” of the uncertainty space, which describes both environmental variability and unknown values of model parameters, can be “learned” from analyzing replicate measurements of the pristine state. The damage detection procedure is based on developing an auto-regressive time series model, comparing its predictions to measurements, and summarizing the results in the form of a scalar-valued damage indicator. We demonstrate that, when it is embedded within the analysis of robustness, this algorithm is capable of separating the various states of the frame structure. A component essential to the success of this technology is the ability to explore model parameters while preserving their correlation structure. This is achieved by using a correlation-preserving sampling scheme based on the principal component decomposition of the correlation.

Our hypothesis, which is that examining how the uncertainty space changes in time might lead to superior diagnostics of structural damage as compared to only monitoring a damage index, is confirmed by these preliminary results. While this methodology holds promise for robust SHM, future research efforts should consider incorporating more of the statistical information that, while it is generated in our analysis, was not fully exploited. Other potential improvements would be to, first, use or develop more sophisticated indicators of damage and, second, generalize the procedure to multiple-input, multiple-output time series modeling.

Acknowledgments The first three authors acknowledge the support of the Engineering Institute and 2015 Dynamics Summer School at the Los Alamos National Laboratory (LANL). The fourth author acknowledges support from the Advanced Scientific Computing program at LANL. The fifth author is grateful for continued support received from the Advanced Certification Campaign at LANL. LANL is operated by the Los Alamos National Security, L.L.C., for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396.

References

1. Sigurdardottir, D.H., Glisic, B.: On-site validation of fiber-optic methods for structural health monitoring: streicker bridge. *J. Civil Struct. Heal. Monit.* **5**(4), 529–549 (2015)
2. Yarnold, M.T., Moon, F.L.: Temperature-based structural health monitoring baseline for long-span bridges. *Eng. Struct.* **86**, 157–167 (2015)
3. Figueiredo, E., Park, G., Figueiras, J., Farrar, C.R., Worden, K.: Structural health monitoring algorithm comparisons using standard data sets. Technical Report LA-14393, Los Alamos National Laboratory, Los Alamos (2009)
4. Fugate, M., Sohn, H., Farrar, C.R.: Vibration-based damage detection using statistical process control. *Mech. Syst. Sign. Process.* **15**, 707–721 (2001)
5. Stull, C.J., Hemez, F.M., Farrar, C.R.: On assessing the robustness of structural health monitoring technologies. *Int. J. Struct. Heal. Monit.* **11**(6), 712–723 (2012)
6. Ben-Haim, Y.: *Info-Gap Decision Theory: Decisions Under Severe Uncertainty*, 2nd edn. Academic Press, Oxford (2006)
7. Farrar, C.R., Worden, K.: *Structural Health Monitoring: A Machine Learning Perspective*. Wiley, New York (2012)
8. Worden, K., Manson, G.: Visualisation and dimension reduction of high-dimensional data for damage detection. 17th International Modal Analysis Conference, Orlando (1999)

9. McKay, M.D., Beckman, R.J., Conover, W.J.: A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics* **21**(2), 239–245 (1979)
10. Van Buren, K.L., Hemez, F.M.: Robust decision-making applied to the NASA multidisciplinary uncertainty quantification challenge problem. *AIAA J. Aerospace Inform. Syst.* **12**(1), 35–48 (2015)
11. Hemez, F.M., Van Buren, K.L.: Designing a mechanical latch for robust performance. 34th International Modal Analysis Conference, Orlando (2015)