A Comparative Study of Approaches to Damage Detection

R.J. Barthorpe[1], K.Worden[1], C.Surace[2], & G. Demarie[2]

[1]*Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, UK*
E-mail: *r.j.barthorpe@sheffield.ac.uk, k.worden@sheffield.ac.uk*

[2]*Department of Structural and Geotechnical Engineering, Politecnico di Torino, Italy*
E-mail: *cecilia.surace@polito.it*

Keywords: Structural Health Monitoring, damage detection, diagnostic methodologies.

SUMMARY. The objective of the current paper is to discuss the relative strengths and weaknesses of the two main approaches to the diagnostic element of Structural Health Monitoring (SHM) and to try to indicate where they are best used in practice. Given the degree of commonality that is apparent, the paper will also discuss how the two approaches can support each other in the development of best practice and some speculation will be made as to how the approaches might be combined in order to exploit the strengths of both.

1 INTRODUCTION

Current approaches to the diagnostic problem which is central to the field of Structural Health Monitoring (SHM) are usually based on two main possibilities: an inverse problem formulation and a machine learning approach. The first of these approaches, often called the model-based approach is usually applied by constructing a physics-based model of the structure of interest (e.g. a Finite Element (FE) model) and correlating it with experimental data. Once the model is established, it can be used in a monitoring phase by periodically updating the parameters of the model, usually by linear-algebraic methods. The nature of the problem means that the linear-algebraic formulation is often ill-posed and requires careful regularisation [1,2]. The alternative approach to diagnostics in SHM, often called the data-based approach, also involves the construction of a model, but this model is usually statistical. The model is established by means of machine learning or pattern recognition and may involve the use of classifiers or novelty (outlier) detectors [3,4]. It must be recognised that the problem of implementing a credible SHM strategy in any real-world context is much more wide-ranging than the choice of a diagnostic methodology. The broader aspects of SHM are however, not discussed here; the reader may consult [5] and [6] for more background; reference [7] must be considered as the current definitive guide to the subject.

Both of the approaches discussed above have substantial support in the literature of SHM; however, they arguably have different strengths and weaknesses, which potentially make their domains of application problem-dependent. The methods also show a degree of commonality which is sometimes overlooked. In the first case, as observed above, both approaches can be said to be model-based. The distinction is in the type of model. If one classifies models into white, grey and black-box models according to their degree of *a priori* physical content; one would observe that the inverse problem approach seeks to establish a white-box model, while the machine learning approach uses a grey or black-box model. The advantage of the former is precisely that it exploits any available physical knowledge of the system of interest; the advantage of the latter is that it automatically accommodates any uncertainty in the specification of the system or structure. Another common aspect of the approaches is that they require measured features from the structure. In the inverse problem approach, features are generally needed in order to update the *a priori* physical model in order to bring the normal condition model into better accord with reality. Features are also required for the damage identification task. These may or may not be the same features that were employed for updating the normal condition model. The machine learning approach uses measured features in order to form the statistical model and to reduce the dimensionality of the problem as far as possible. In both cases, the selected features must be sensitive to the damage. For the inverse problem this is a requirement for a non-trivial update; for the machine learning approach, the features are essentially everything and must reflect any information about the damage. Feature selection, is therefore an issue for all SHM methodologies, but is arguably most discussed in the machine learning context.

2 CLASSIFICATION OF MODELS

The discussion later in the paper will require the specification of a taxonomy which allows one to distinguish between classes of models. The two most common means of distinguishing model types are covered by the following.

2.1 *Data-driven and law-driven models*

A good reference on this classification of models is [8]. A convenient way of expressing the differences between the two types of model is by means of ‘bullet points’.

Law-driven models

- Based on accepted laws attributed to the system – ‘physical’ in nature
- Suited to prediction, potentially for unobserved system states.
- May be used to inform critical data acquisition decisions
- Typically highly-parameterised
- Generally do not accommodate uncertainty.

Data-driven models

- Based on observed input/output relationships – ‘statistical’ in nature.
- Suited to recovering inputs from observed outputs.
- May be parsimoniously parameterised.
- Naturally accommodate uncertainty.

This is a useful picture; however, one should regard these statements as representing ‘extreme’ viewpoints; things are seldom black and white. Consider the suitability of the two paradigms for ‘prediction’ purposes. It is well-known that data–driven models like neural networks should only be used in situations where the input data do not depart dramatically from those used during the ‘training’ of the model [9]; in other words, such models can only be used to *interpolate* with any real confidence. In contrast, one would imagine that an appropriately derived law-driven or physics-based model could be used to *extrapolate* i.e. make predictions in situations removed from the current SHM context; physics should after all be universal. However, one should bear in mind that law-driven models may be the best that physics has to offer, yet still be subject to restriction; for example, the Navier-Stokes equation is restricted to situations where one has confidence in a continuum assumption. Further, when one describes the Navier-Stokes equation as a physics-based model one is misrepresenting the situation in terms of prediction capability as the equation has no known analytical solutions; prediction is only possible by passing to numerical solution methods which are themselves subject to limitations and represent a different class of model to the original partial differential equation. Finally, one can observe that a perfect specification of a deterministic model still does not guarantee prediction accuracy; this is one of the hard lessons learned from the discovery of deterministic chaos [10]. One can also consider the question of ‘parameterisation’ of the models. Although it is stated above that law-driven models may be ‘highly-parameterised’; this need not be the case. If one considers a large FE model, there may be very many parameters indeed; in principle, each individual element could have independent material parameters. In fact, only a small number of ‘substructure’ material constants will be needed. Of the huge number of coordinates which specify the mesh geometry; only the subset which fixes the geometry of the structure of interest matter in any real sense. ‘Internal’ nodal coordinates can be varied more or less with impunity (as long as one does not violate aspect ratio constraints etc.). It is this fact that only a very small subset of the parameters ‘matters’ which allows FE model updating to be computationally feasible for model improvement. For SHM the additional requirement is that the model should be capable of representing the structure in all damage states of interest. Meeting this requirement while avoiding an explosion in the number of required parameters (and subsequent issues of ill-conditioning) is perhaps one of the greater challenges faced in the application of model updating based methods.

The parallel statement that data-driven models can be ‘parsimoniously-parameterised’ is by no means a general rule either; neural networks may need hundreds of internal parameters in order to capture the input-output behaviour of a system with appropriate fidelity; this observation impacts considerably on the applicability of data-driven models, as it imposes severe constraints on the amount of data which must be available for ‘training’.

2.2 *Black, white and grey box models*

"Here they come, every colour of the rainbow: black, white, brown" [11].

It is informative to consider the nature of the information upon which a model is based. A white box model is built solely upon the best possible understanding of the underlying physics of the system. For a purely white model, no assumptions or approximations regarding model structure would be made, and the ‘whiteness’ of a given model reflects the depth to which a complete physical understanding is pursued. In practice, limits on our understanding of the physical universe preclude the existence of pure white box models. The discussion on the Navier-Stokes equation above makes this point clear. One should also note that both the Navier-Stokes equation itself and the discretisation usually necessitated for solution are often brought together under the term ‘white box.

Conversely, a black box model seeks to describe the behaviour of a system with no reference to its internal structure or ‘physics’. Such models are instead built solely around observed input-output behaviour. This does not always present practical limitations as the universal approximation capability of many machine learning model paradigms means that a black-box model can in principal capture input-output behaviour perfectly; however, as discussed above, for predictions with such models it is not usually safe to stray too far away from the situations in which the training data were generated.

A grey-box model is a compromise between these two extremes i.e. a model for which the physics dictating the input-output relationships are partially understood *a priori*, but which allows for the inclusion of approximations from empirical observations. The relative 'whiteness' or 'blackness' of the model may be viewed as dependent upon the number and quality of the assumptions made in specifying the physical understanding of the system, and the degree to which the model relies upon approximations made from the observed data. A good example here may be in specifying an initially linear FE model for a structure. If experiment were to make clear that the actual structure was in fact nonlinear, one could add individual nonlinear elements and calibrate their coefficients using experimental observations. Even if the nonlinearities are non-polynomial in nature, they will always be approximated effectively by an appropriate high-order polynomial; such a 'model' nonlinearity is therefore non-physical and converts the initial white FE model to a grey box.

3 THE MODELLING TASK FOR DAMAGE IDENTIFICATION

The purpose of this section is to discuss the modelling capabilities of the two main approaches: law-based and data-based. The 'true' structure may be regarded as a function (or functional) of its inputs. This function specifies the mapping of inputs to responses for all states of the structure, both damaged and undamaged. For damage identification, the modelling task is to specify a model capable of approximating the true structure across all states of interest.

The discussion is illustrated through consideration of a space of functions or functionals encompassing all possible model structures. The 'true' structure is denoted as *T* in the following figures. Where model updating has been employed to reduce the residual distance between model and structure, the 'optimal' model is denoted as *O* in the figures. In the interests of visualisation, the abstract concept of an infinite-dimensional functional space is presented in two dimensions.

3.1 *Physics-based modelling*

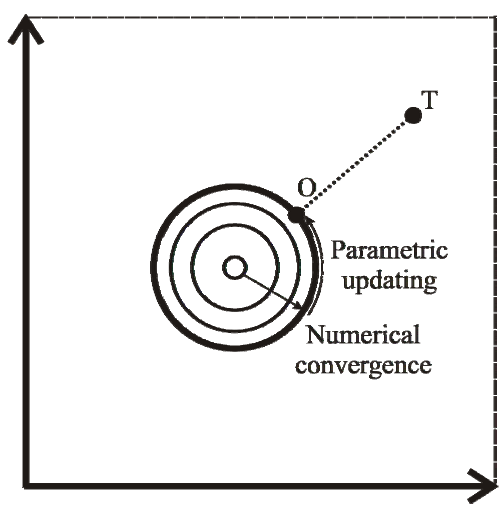

Figure 1: Functional space portrait of a linear physics-based model.

Suppose one begins by considering the family of linear physics-based models as depicted in Figure 1. This class of models, by nature of its restrictions, forms a subset or spans a subspace of all the possible functional forms for a model. The particular subspace spanned is dictated by the model form selected and the values assigned to its parameters. The extent of this subspace (the modelling 'power' or capability) grows with the dimension of the model – as more degrees of freedom are added to, say, a finite element discretisation of the structure, so the number of possible models expands. One can specify a complexity parameter: in this case the number of elements is meaningful. So for linear FE models, one has a complexity parameter: $N_{phys} = \{N_{elements}\}$.

The creation of a numerical model necessarily involves the discretisation of continuous physical laws, manifested as specification of time-steps or element sizes. The discretisation error arising from this process limits the capability of the model to accurately reflect the physics of the true system. As the step- or element-size is decreased (and thus the order of the model increased) convergence of the subspace to some outer limit or boundary would be expected.

Within the spanned subspace, the 'optimal' model is that which minimises the distance (in terms of the geometry of the function space) between the model and that of the 'true' structure. This optimisation may be guided by direct measurement of parametric values and/or calibration of parameters through updating. Where parametric updating methods are employed for this purpose, the optimisation process will typically be constrained in order to ensure that the model parameters remain physically representative. e.g. for FE updating, a principled approach to choosing the 'closest' model within the class could be based on least-squares minimisation of residual errors; possible constraints on the minimisation might be the requirement to preserve the sparsity structure of the physical matrices during updating [12]. If the true system falls within the explanatory range of the model class, the updating procedure will bring the distance down to zero. If the true system is without the model class as depicted in Figure 1, one can only hope to get as close as possible.

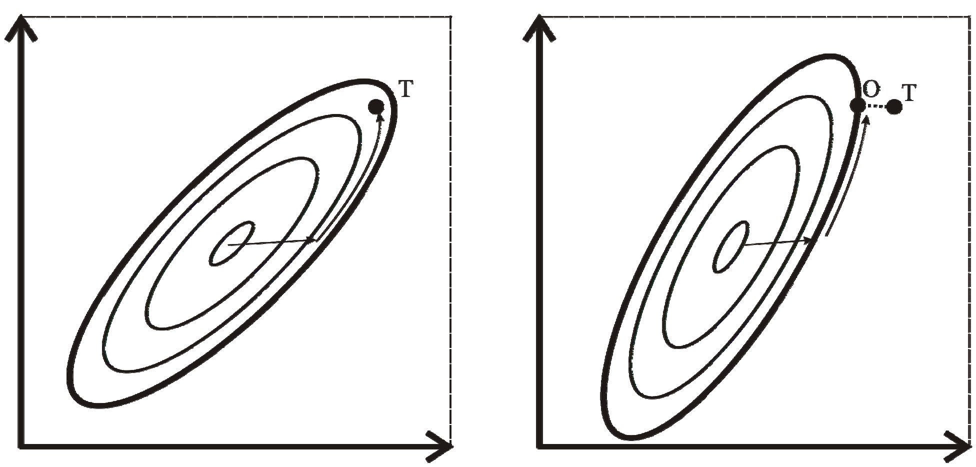

Figure 2: Functional space portrait of a non-linear physics-based model where:
a) the included non-linearities accurately reflect those of the true structure, and
b) there are discrepancies between the included non-linearities and those of the true structure

One can now progress to, say, nonlinear FE models. Such models will have at least the explanatory power of linear FE models, but will clearly span a greater volume of the function space. Staying within the dictates of law-based modelling, adding finite numbers of specific nonlinearities will mean the linear FE class is only extended in certain directions in the function space. In general, this may well mean that the extended class still does not include the true system (as in Figure 2b); however, it may be that it does (Figure 2a). The complexity parameters for the class must also specify the number of 'nonlinearities' added: $N_{phys} = \{N_{elements}, N_{non\text{-}linearities}\}$. Any given physics-based class, e.g. nonlinear FE models, will still be subject to restrictions on the explanatory power even if the number of elements grows without bound. This example represents the case discussed earlier where addition of extra terms increases the explanatory power of a white box model but converts it into a grey box in the process. Concentrating on a linear FE model with 'added' nonlinearity is extremely relevant for SHM as many damage types will convert an initially linear structure into a nonlinear one. Further, the exact form of the nonlinearity relevant to a fatigue crack in a metal or a delamination in a composite laminate may not be precisely known.

3.2 *Data-based modelling*

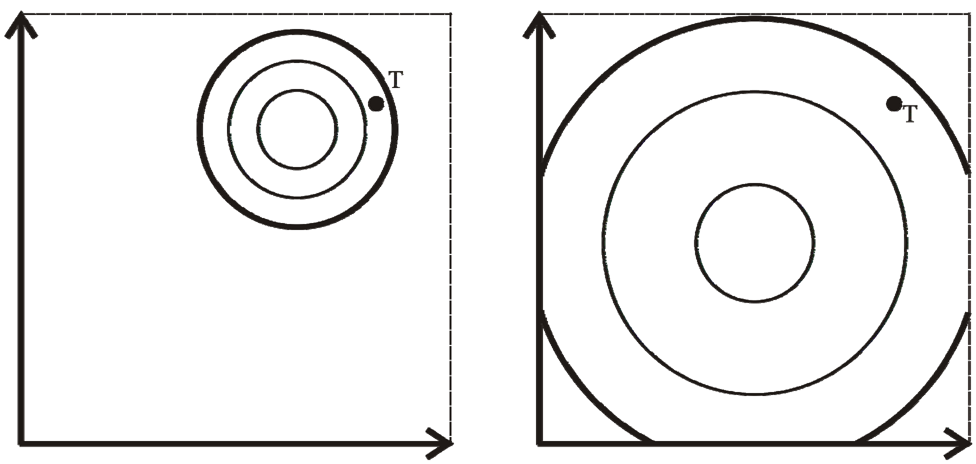

Figure 3: Functional space portrait of a data-based approach where:
a) the structure of the data-based model is 'close' to the true structure, and
b) the structure of the data-based model is not 'close' to the true structure.

As in the case of physics-based models, one can regard a given data-based paradigm as spanning a subspace of possible functions/functionals. However, there is a critical difference. Many classes of data-based models can be proved capable of acting as universal approximators; these classes include: multi-layer perceptron neural networks, radial-basis function networks [9] and Support Vector Machines [13]. The universal approximation capability means that as the complexity parameter (number of free weights etc.) increases without bound, any function/functional can be represented arbitrarily accurately. In terms of the figures here, this means that as the number of parameters increases, any point in the space is reachable and one should always encompass the 'true' point T for some complexity parameter: $N_{data} = \{N_{free\ params}\}$. Unfortunately, the greater the number of parameters, the greater the amount of training data that is required and this is the critical problem in the context of SHM. The 'true' function in an SHM context is often a classifier which maps data to a class label specifying the state of health of a given structure. If the classifier is a dual class (healthy/damaged) novelty detector, it may only need to be trained on examples of data from the healthy structure; this is called unsupervised learning. If the classifier is more refined and is required to indicate location or severity of damage, training data from the damage states will be needed; the problem is now one of supervised learning. In an engineering context, this presents a great problem as it may not be economically feasible to generate fault data exemplars from high-value structures like aircraft or bridges. Training of a complex model (many parameters) may demand a great deal of training data in a situation where acquiring any damage state data at all will be a formidable problem

3.3 *Hybrid approach*

The 'hybrid' approach presents a potential solution to the problem just discussed. One can use a physics-based approach to establish a model which explains as much as possible of the function/functional behaviour as possible. The situation, as depicted in Figure 4a, will be that the initial physics-based model will establish a point O, close to the truth. From this point, a data-based adjunct will be able to bridge the gap between O and T because of the universal approximation capability of the data-based approach. The advantage here is that the data-based component will potentially explain the 'residual' between O and T with a need for fewer parameters (than a full data-based model) and will therefore require less training data.

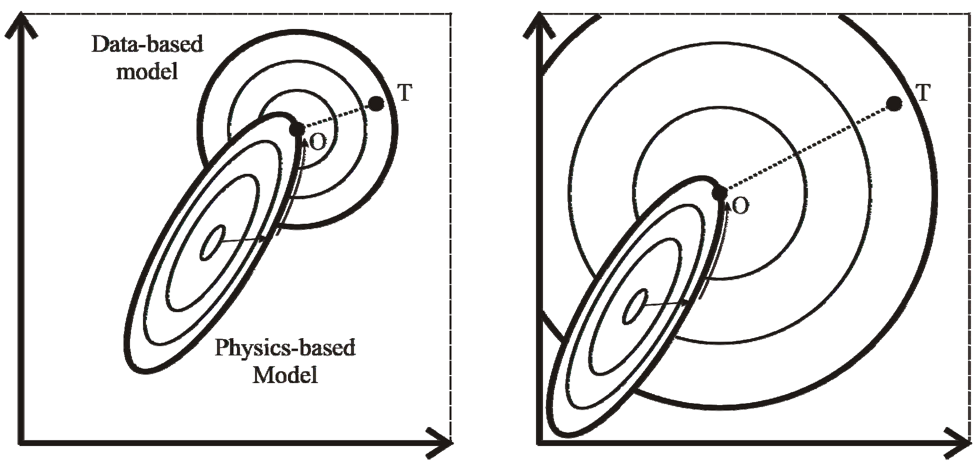

Figure 4: Functional space portrait of a hybrid physics-based and data-based approach where
a) the physics-based model is 'close' to the true structure, and
b) the physics-based model is not 'close' to the true structure.

The complexity measures in this case are $N_{phys} = \{N_{elements}, N_{non\text{-}linearities}\}$, $N_{data} = \{N_{free\ params}\}$. In the situation shown in Figure 4b, the initial physics-based model has less explanatory power than that in Figure 4a and as a result the parameter count for the data-based component will be higher.

4 PRACTICAL CONSIDERATIONS

Each method may be considered as comprising a training phase and a monitoring phase. The training phase for the data-based approach is precisely as described in all the relevant texts in machine learning or pattern recognition. The monitoring phase comprises the application of the learnt ‘rules’ to newly presented data from the possibly damaged structure. In terms of the law-based approach, the ‘training phase’ means the initial building of the physics-based model on the basis of all available prior information followed by updating and validation exercises using experimental observations to bring the model further into accord with reality. The monitoring phase for a physics-based approach will typically comprise further updating steps on the basis of newly presented data from the structure.

For the hybrid approach the training phase is considered in the discussion below to involve the development of the ‘optimal’ law-based model, and the training of a data-based model using the (bias-corrected) predictions of this law-based model. The monitoring phase is once again the application of the data-based model to new data. It should be noted that the objectives of the law-based modelling task are thus somewhat different for the hybrid approach and the purely physics-based approach. For the hybrid approach it is advantageous for the reliance on deterministic updating to be superseded by application of the rather broader concept of probabilistic model validation. The broad aim of probabilistic model validation is the pursuit of of quantified levels of confidence in prediction [14].

4.1 *Features employed*

As observed, both approaches have a training phase which requires the availability of experimental data. The exact type of data used will have a critical effect on the construction of the diagnostic. The most important point is that the data, expressed through multivariate data vectors called ‘features’, must be sensitive to damage.

For physics-based methods in general, the selection of the feature type is heavily dependent upon being able to identify the feature both from the model and the experimental data. In the model-updating literature, this has historically led to the broad adoption of modal properties (primarily frequencies and modeshapes) or FRFs. Comparison between experiment and the initial model leads to residuals that must be minimised when updating structural models. Often, the updating step will require analytical results, specific to given features e.g. sensitivity-based updating requires the formulation of derivatives of the error function with respect to the features; these derivatives have been computed and are available in the literature for the standard modal features [12]. This reliance upon modal properties has been maintained despite the insensitivity of modal responses to damage being well documented. It should be noted that a third category of modal characteristic, damping, is rarely considered for forming residuals for model improvement or as a feature for damage detection, due to difficulties associated with characterising and measuring damping.

Despite the effort invested in reducing the residual between the responses of a physics-based model and the ‘true’ structure, some degree of offset will inevitably remain. This offset arises in part from the impossibility of achieving a true white-box level understanding of the physics of the structure. Bias correction techniques may be used in an effort to circumvent this error.

The restriction to features that can be identified and correlated between model and structure is removed in the purely data-driven case. For the machine learning approach, ‘restrictions’ on the features are only to what can be measured from the structure: modal characteristics, time-domain data, spectral data, strain histories, images. Given this lack of restriction on the type of feature, emphasis is instead placed upon reducing the dimensionality of the problem as far as is possible (and thus reducing the requirement for infeasible amounts of training data) and selecting a feature set that is robustly indicative of damage.

The degree of ‘restriction’ on the feature set falls somewhere between these two levels for the hybrid approach. The decision on whether a particular response of the model should be considered a candidate feature will largely be dictated by whether it can be predicted with a satisfactory degree of confidence.

4.2 *Feature selection*

It is sensible to make a distinction between feature specification and feature selection. The broad class of features available will be determined by the choice of sensors and their number and distribution; as such this must be considered at the operational evaluation stage as discussed in [5]. A great advantage of employing a physics-based approach is that the model may be used to guide the sensor placement and feature selection tasks, and it would in principle be possible to assess the sensitivity of each individual feature within the feature set to the presence of damage. Examples of this are, however, relatively rare in the published literature related to model-updating based SHM, although Fritzen et al. [15] have employed parameter set reduction techniques. In the inverse model-based approach, as discussed above, the candidate feature set used is typically specified *a priori*, and is often the same for both the model-improvement stage and the structural monitoring stage. This is despite the objectives of these two stages being arguably somewhat different. The feature set used is defined largely by convention and convenience, rather than through analysis of suitability for damage detection.

In direct contrast, feature selection – the reduction of a candidate set to a maximally sensitive set - is very widely discussed in the machine learning context, where the features are essentially everything. A practical consideration here (returned to below) is that data is unlikely to be available for all damage states, placing a restriction on the domain for which the features are assessed. A solution making full use of the hybrid formulation may be to use the physic-based model to provide probabilistic representations of a candidate feature set, and to apply feature selection methods developed in the machine learning context to this set.

4.3 *Treatment of Uncertainty/Variability*

There are many potential sources of uncertainty in the SHM task. Several of these, including boundary condition variability and operational and environmental effects are common to both approaches. The use of physics-driven models introduces an additional set of model-form and parametric uncertainties. The two approaches handle uncertainties in markedly differing ways.

There are two elements to consider when assessing the importance of uncertain factors:

- are the features that are being observed sensitive to the uncertain factors?
- What degree of uncertainty can be ascribed to the uncertain factors?

An advantage of the machine learning approach is that it can automatically accommodate any uncertainty in the specification of the system or structure. This is, however, contingent upon the learning algorithm being presented with training data that is representative of the variabilities and uncertainties that are expected in practice. A common, yet non-trivial, example is that of environmental effects. It has been observed in numerous studies that temperature effects can have a significant effect upon the features employed for damage identification. In some cases, it has been found that the observed features display greater sensitivity to environmental effects than to the damage of interest [16]. The pitfalls of training a classifier using data gathered at a single temperature point in such a scenario are immediately apparent, and it would be expected that a classifier trained at one temperature point may perform poorly when extrapolating to changing environmental conditions. In this scenario, the possibility may exist to treat the variable factor (temperature) as observable, and to incorporate it into the feature vector presented to the machine learning algorithm. Where the variable factor is not readily observable, as may be the case for boundary conditions, a different approach may be taken. Effort may be directed towards recording training data that captures the range of variability expected in practice, for example through employing blocking and randomisation as appropriate in the test programme. An initial analysis of the sensitivity of candidate features to particular uncertain factors may aid this process.

In the inverse-problem approach to SHM, the rule has very largely been to build 'crisp' models which do not accommodate variation; departures from this rule have recently emerged via the use of stochastic FE [17] and fuzzy FE [18] etc. The quantification and propagation of uncertainty is in contrast a core activity in probabilistic model validation approaches. Sensitivity analysis is conducted using the initially-developed model to identify those factors to which model responses are sensitive. The experimental and modelling effort is directed at characterising the uncertainty in these factors, and the subsequent variability in the model outputs, The resulting probabilistic predictions may be used to supplement the available experimental data for machine learning in the hybrid approach, Even when uncertainty is explicitly handled in a law-based model; the form of the uncertainty must be guided or fixed on the basis of observational evidence; this may well change a white box model into a grey box in the same way that addition of nonlinear elements would

4.4 .*Practical considerations*

The discussion so far has been largely related to the advantages and disadvantages of models in terms of their explanatory power, their ability to encode uncertainty and other matters which have been discussed in largely abstract terms. The intention with this section is to look at practical issues including those of resourcing.

In terms of cost (time and money), physics-based models are undoubtedly time-consuming and challenging to build, develop and validate. They rely on the existence of expert and talented model builders with considerable training, both in the generalities of model building and the specifics of individual software packages. The software packages themselves may be costly in terms of initial licensing and maintenance. Large models will require intensive computing facilities with the associated operation and maintenance overheads. Robustness under uncertainty can be an issue for law-based models; these are typically calibrated at a single design point (the normal condition), and validated for responses not used for updating, rather than for other input states. Robustness and sensitivity are often left unassessed. Also, additional epistemic uncertainties will generally be introduced through modelling choices and simplifications. Law-based models are typically over-parameterised; this makes updating a typically ill-conditioned task. (In order to improve the stability of solution, regularisation may be employed to reduce the dimension of the parameter space [14]). Finally, during the monitoring phase, any assessment of integrity will require a full model update step. There are of course many points in favour of the law-based approach. First, law-based models can potentially be used to extrapolate; if the actual physical laws underlying the model extend beyond the initial context, then the constructed model will also extend. Secondly, the actual computational cost of the update step may not be excessive. A huge advantage of the law-based model is concerned with the possibility of observing multi-site damage. Consider the situation where the structure under investigation is a cantilever beam with a potential fatigue crack.

Once the undamaged structural model has been validated, and the crack model has also been validated, there is no further problem with modelling the multiple crack scenario; this is not at all the case for the data-based approaches.

In discussing the disadvantages of the data-based approach, one must begin with the 'elephant in the room'. Sourcing data from the structure in its damaged state is unlikely to be feasible in scenarios such as large-scale structures. One clearly cannot conceive of physically damaging an aircraft in multiple ways in order to accumulate data for supervised learning. Even if multiple damage cases become available, there is a potential requirement to account for statistical variation in the data. Finally, even is it is possible to acquire some data from damaged structures it will certainly be unfeasible to acquire data corresponding to multi-site damage, where the number of damage states for which data would be required grows factorially with the number of discrete damage sites identified. This is the main problem faced by data-based SHM. If only damage detection is required, things become much more positive as only measured data from an undamaged structure will be required. A related problem concerns the amount of training data required. If the machine learning model structure has many adjustable parameters, it may demand an unfeasible amount of training data – data, even from undamaged structures, does not come without a cost. In terms of model development; this can be accomplished using machine learning software which is comparatively quick to master (compared with a nonlinear dynamic FE solver for example). Training the model may well be time-consuming, but then runs will be extremely fast during the monitoring phase.

5 CONCLUSIONS

As the paper is essentially a discussion document, there is no real need for detailed conclusions. The paper simply discusses the differences and commonalities between the law-based and data-based approaches to structural health monitoring. The main issues discussed here relate to: the relative explanatory power of types of models, the accommodation of uncertainty in models and finally, practical issues in implementing diagnostic strategies. While these issues surely do not exhaust issues for possible discussion, the authors feel that their importance justifies their prominence here. It is observed that the two approaches discussed have their own individual strengths and weaknesses and that a hybrid approach may be possible that exploits the capabilities of both. The development of such a hybrid approach is the subject of ongoing research

References

[1] Friswell, M.I., "Damage identification using inverse methods," in *Dynamic Methods for Damage Detection* (CISM Courses and Lectures vol. 499, eds. A. Morassi and F. Vestroni), Springer Wien, New York, 13-66 (2008).

[2] Worden, K. and Friswell, M.I., "Modal vibration-based damage identification," in *Encyclopedia of Structural Health Monitoring*, eds. C. Boller, F.-K. Chang and Y. Fujino, John Wiley and Sons (2009).

[3] Worden, K. and Dulieu-Barton, J.M., "An overview of intelligent fault detection in systems and structures," *Structural Health Monitoring, an International Journal*, **3**, 85-98, (2004)..

[4] Worden, K. and Manson, G., "Applications of machine learning to structural health monitoring," *Philosophical Transactions of the Royal Society, Series A*, **365**, 515-537, (2007).

[5] Farrar, C.R. and Worden, K., "An introduction to structural health monitoring," *Philosophical Transactions of the Royal Society, Series A*, **365**, 303-315, (2007).

[6] Worden, K., Farrar, C.R. and Dulieu-Barton, J.M., "Principles of structural degradation monitoring," in *Encyclopedia of Structural Health Monitoring*, eds. C. Boller, F.-K. Chang and Y. Fujino, John Wiley and Sons (2009).

[7] Boller, C., Chang, F.-K. and Fujino, Y., (Eds.), *Encyclopedia of Structural Health Monitoring*. John Wiley and Sons (2009).

[8] Saltelli, A. and Ratto, M., *Global Sensitivity Analysis: the Primer*. John Wiley and Sons, Ltd, Chichester, (2008).

[9] Bishop, C.M., *Neural Networks for Pattern Recognition*. Oxford University Press, (1998).

[10] Schuster, H.-G., *Deterministic Chaos*. John Wiley and Sons Australia Ltd., (1998).

[11] Colemanballs. http://en.wikipedia.org/wiki/Colemanballs. Accessed 11/08/2009.

[12] Friswell, M.I. and Mottershead, J.E., *Finite Element Updating in Structural Dynamics*. Kluwer Academic Publishers, (1995).

[13] Christianini, N. and Shawe-Taylor, J., *An Introduction to Support Vector Machines and other Kernel-based Learning Methods*. Cambridge University Press, (2000).

[14] Hemez, F.M., "Uncertainty quantification and the verification and validation of computational models," In *Damage Prognosis for Aerospace, Civil and Mechanical Systems*, D.J. Inman, C.R. Farrar CR, V. Lopes Jr, V. Steffen Jr, (Eds), John Wiley and Sons Ltd, London , (2004).

[15] Fritzen, C.-P., Jennewein, D. and Kiefer, T., "Damage detection based on model updating methods," *Mechanical Systems and Signal Processing*, **12**, 163-85, (1998).

[16] Sohn, H., "Effects of environmental and operational variability on structural health monitoring," *Philosophical Transactions of the Royal Society, Series A*, **365**, 539-560, (2007).

[17] Ghanem, R.G. and Spanos, P.D., *Stochastic Finite Elements: A Spectral Approach*. Dover Publications Inc., (2003).

[18] Lallemand, B., Cherki, A. Tison, T. and Level, P., "Fuzzy modal finite element analysis of structures with imprecise material Properties," Journal of Sound and Vibration, 220, 353-365, (1999).