

Model validation: Correlation for updating

D J EWINS

Department of Mechanical Engineering, Imperial College of Science,
Technology and Medicine, London, UK
e-mail: d.ewins@ic.ac.uk

Abstract. In this paper, a review is presented of the various methods which are available for the purpose of performing a systematic comparison and correlation between two sets of vibration data. In the present case, the application of interest is in conducting this correlation process as a prelude to model correction or updating activity.

Keywords. Correlation; model validation; model updating; mode shape comparison; modal assurance criterion; frequency response assurance criterion.

1. Introduction

In this paper, we turn our attention to the role of using modal test data for the purpose of refining the theoretical model which will be used for the design optimisation process. There are many different names given to the tasks involved in this refinement. ‘Model updating’ is a common one, as is ‘model validation’ and ‘reconciliation of test and analysis’. Here, and in the following paper, we shall review the range of activities which span the relatively user-intensive techniques of manual updating through to the more automatic algorithms being developed today. We shall start by exploring the process of test/analysis comparison and correlation as this is an essential first step on the way to updating the theoretical model. If we are unable to obtain a satisfactory degree of correlation between the initial theoretical model and the test data, then it is extremely unlikely that any form of model updating (correcting the model to match the test data) will succeed. Thus, a successful correlation is crucial for the success of model updating.

2. Comparison of experiment and prediction

2.1 *Different methods of comparison*

Probably the single most popular application of modal testing is to provide a direct comparison between predictions for the dynamic behaviour of a structure and those actually observed in practice. Sometimes this process is referred to as ‘validating’ a theoretical model, although to do this effectively several steps must be taken. The first of these steps is to make a direct and objective comparison of specific dynamic properties, measured versus

predicted. The second (or, perhaps still part of the first) is to quantify the extent of the differences (or similarities) between the two sets of data. Then, the third step is to make adjustments or modifications to one or other set of results in order to bring them closer into line with each other. When this is achieved, the theoretical model can be said to have been *validated* and is then fit to be used for further analysis. In this section we shall be concerned with the first (and, to some extent, the second) of these stages, dealing with the others later.

In most cases, a great deal of effort and expense goes into the processes which lead to the production of an experimentally-derived model on the one hand (subsequently referred to as the 'experimental' model or data) and a theoretically-derived (or 'predicted') model on the other. This being so, it is appropriate to make as many different types or levels of comparison between the two sets of data as possible. As discussed earlier, we can identify three types of dynamic models, loosely called 'Spatial', 'Modal' and 'Response'. It is now convenient to return to this classification and to try to make comparisons between experiment and prediction at each (or at least more than one) of these. Thus we shall discuss comparisons of response characteristics and of modal properties, as both of these provide many opportunities for useful correlation between experiment and theory. Comparisons of spatial properties are more difficult, however, and we shall not discuss this aspect now.

Whichever medium is used for comparison purposes, either one or the other model will have to be developed fairly extensively from its original form and what is the most convenient format for one case will be the least accessible for the other. This situation derives from the different routes taken by theoretical and experimental approaches to structural vibration analysis. However, in closing these general remarks, it is appropriate to reinforce the recommendation to make as many different types of comparisons as possible and not just to rely on one, usually the first one that comes to hand, or mind.

One further comment, and concept, which should be introduced is that of *verifying* the models to be compared. The concept of a verified model is different to that of a validated model in the following respects. A model can be said to be *verified* if it contains the correct features, most importantly the appropriate number and choice of DOFs, to represent the behaviour of the structure; a model is said to be *valid* if the coefficients in that model are such as to provide an acceptable quantitative representation of the actual behaviour. It will be seen that a model can only be validated after it has been verified. This means that we should not embark on lengthy comparison or correlation procedures unless we are first satisfied that the two models to be used are compatible with each other and their intended roles.

2.2 Comparison of modal properties

While there is no compelling reason for choosing one rather than the other, we shall start our comparison procedures with those based on modal data and follow with those which use response properties. Although the response data are those most directly available from test for comparison purposes, some theoretical analysis packages are less than convenient when it comes to predicting FRF plots. This is largely because of the requirement that all (or at least a large proportion) of the modes must be included in the calculation of a response characteristic. By contrast, modal properties can be predicted individually and comparisons can be confined to specific frequencies or to specific frequency ranges with much greater facility for the analyst. However, such a comparison does place additional




Figure 1. Plots of measured vs predicted natural frequencies, order decided by mode number (a) and mode shape (b).

demands on the experimental route as it requires the measured data to have been subjected to a modal analysis or curve-fitting procedure in order to extract the corresponding modal properties from the test. In spite of this requirement, comparisons of modal properties are perhaps the most common and we shall now describe a number of methods which may be employed to that end.

2.2a Comparisons of natural frequencies: The most obvious comparison to make is of the measured versus the predicted natural frequencies. This is often done by a simple tabulation of the two sets of results but a more useful format is to plot the experimental value against the predicted one for each of the modes included in the comparison, as shown in figure 1a. In this way it is possible to see not only the degree of correlation between the two sets of results but also the nature (and possible cause) of any discrepancies which do exist. It is important to stress, however, that the points plotted in this way must be of the measured and predicted natural frequencies of corresponding modes. It is not sufficient simply to plot the first, second, third... measured natural frequencies against the first, second, third... predicted values as there is no guarantee that the first three measured modes correspond one-for-one with their predicted counterparts. Some positive identification of each mode with its counterpart is essential, to provide a set of Correlated Mode Pairs (CMPs), and for this, recourse must usually be made to the mode shapes correlation methods discussed in the next section. An example of such a situation is shown in figures 1a and b, where the first of these plots shows the natural frequencies plotted simply in ascending order, while the second shows them correctly paired (using information about their modes shapes – not shown here) but displaying the poorer level of correlation which actually exists in this case.

Once sorted, the points on this plot should lie on or close to a straight line of slope 1. If they lie close to a line of a different slope then almost certainly the cause of the discrepancy is an erroneous material property used in the predictions. If the points lie scattered widely about the 45° straight line then there is a serious failure of the model in representing the test structure and a fundamental re-evaluation is called for. If the scatter is small and randomly distributed about a 45° line then this may be expected from a normal modelling and measurement process. However, a case of particular interest is where the points deviate slightly from the ideal line but in a systematic rather than a random fashion as this situation suggests that there is a specific characteristic responsible for the deviation and that this cannot simply be attributed to experimental errors.

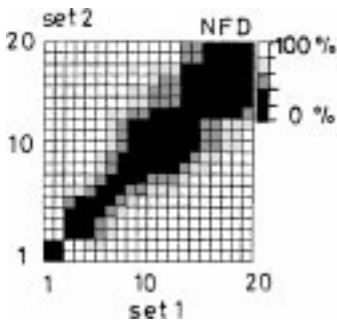


Figure 2. Natural frequency difference (NFD) plot.

There is an inclination to quantify the deviation of the plotted points from the ideal straight line as a means of assessing the quality of the comparison. Although this is indeed useful, it cannot replace the benefit gained from the plot itself as (without employing complicated functions) it is generally insensitive to the randomness or otherwise of the deviations and this is an important feature.

Another possible form of plotting these same data is provided by the Natural Frequency Difference (NFD) table, as shown in figure 2. This is a table which simply plots the natural frequency difference between each possible combination of experimental mode and analytical model mode. This can be used in the automatic selection of correlated mode pairs which is a feature found in more advanced correlation packages.

2.2b Comparison of mode shapes – graphical: When the above procedure is applied in practical cases, it is often more difficult than first anticipated because of the problems of matching the experimental modes with their analytical counterparts. Whereas on simple structures with well-separated modes this pairing generally presents no difficulty, on more complex structures – especially ones with closely-spaced natural frequencies – ensuring that the correlated mode pairs are correctly made becomes more difficult and requires the additional information in each case of the mode shape as well as the natural frequency. Hence it is appropriate to make comparisons of mode shapes at the same time as those of natural frequencies.

In this case, we have rather more data to handle for each mode and one possible way of performing the comparison is by plotting the deformed shape for each model – experimental and predicted – and overlaying one plot on the other. The disadvantage of this approach is that although differences do show up, they are difficult to interpret and often the resulting plots become very confusing because there is so much information included. An alternative plot comprising a single picture which is a display of the mode shape difference can also be difficult to interpret. A more convenient approach is available by making an x - y plot along similar lines to that used for the natural frequencies in which each element in the mode shape vector is plotted, experimental versus predicted, on an x - y plot such as is shown in figure 3. The individual points on this plot relate to specific DOFs on the model and it is expected that they should lie close to a straight line passing through the origin. If, as is often the case, both sets of mode shape data consist of mass-normalised eigenvectors, then the straight line to which the points should be close will have a slope of ± 1 . Once again, the pattern of any deviation from this requirement can indicate quite clearly the cause of the discrepancy: if the points lie close to a straight line of slope significantly different from ± 1 , then either one or other mode shape is not mass-normalised




Figure 3. Plots of measured vs predicted mode shape vectors – (a) single mode; (b) 3 modes; (c) corrected model.

or there is some other form of scaling error in the data. If the points are widely scattered about a line, then there is considerable inaccuracy in one or other set and if the scatter is excessive, it may be the case that the two eigenvectors whose elements are being compared do not relate to the same mode.

This form of presentation has particular value when the deviations of the points from the expected line are systematic in some way, such as is the case in figure 3a. In this event it can be useful to superimpose the plots for several modes so that the basis of the comparison is broadened, and this has been done in figure 3b for the first three modes of the structure. We now see that three of the points on the structure (4, 5 and 6) systematically produce a poor correlation between experiment and prediction although we are not yet in a position to identify which set is in error. From figure 3b it is clear that most of the points for coordinates 4, 5 and 6 do in fact lie close to a straight line but one with a slope considerably different from 45°. If the discrepancy is due to poor analytical modelling (the natural assumption of the experimentalist!), then it might reasonably be expected to differ in extent from one mode to the next. However, this is not the case here and it can be seen that the deviations are consistent with the result which would follow from an incorrect scaling factor on the measured FRF plots pertaining to points 4, 5 and 6 (since all modes would be equally affected by such an error). A repeat of the measurement (and modal analysis) phase in this case, together with the inclusion of some additional coordinates, resulted in the revised plot shown in figure 3c, clearly a much more satisfactory comparison and one achieved using the original analytical model.

At this juncture, it must be observed that the above assumes implicitly that the mode shapes in both cases are real (as opposed to complex) and while it is highly likely that the results from a theoretical analysis will indeed comply with this assumption, those from an experimental source will, in general, not be so simple. Although it is possible to envisage a complex version of the type of plot discussed above by using a third axis to display the imaginary part of the complex eigenvector elements, this is not recommended as it tends to disguise the essential conflict which is inherent in comparing complex (experimental) data with real (predicted) values. It is necessary to make a conscious decision on how to handle this particular problem and that usually adopted is to ‘whitewash’ the measured data by taking the magnitude of each eigenvector element together with a + or – sign depending on the proximity of the phase angle to 0° or 180°. In many cases this is adequate but it is not satisfactory for highly complex modes. No form

of direct comparison between these modes and the real data produced by a typically undamped theoretical model will be effective. In such cases, it becomes necessary to employ one of the complex-to-real transformation or ‘realisation’ procedures which can be found in the literature.

2.2c Comparison of mode shapes – numerical: Several workers have developed techniques for quantifying the comparison between measured and predicted mode shapes (in fact, these methods are useful for all sorts of comparisons, not just experiment vs theory, and can be used for comparing any pair of vectors). As an alternative to the above graphical approach, we can compute some simple statistical properties for a pair of modes under scrutiny. The formulae given below assume that the mode shape data may be complex, and are based on a comparison between an experimentally measured mode shape, $\{\psi_X\}$, and a theoretically predicted one, $\{\psi_A\}$.

The first formula is for a quantity sometimes referred to as the ‘Modal Scale Factor’ (MSF) and it represents the ‘slope’ of the best straight line through the points as plotted in figure 3. This quantity is defined as

$$\text{MSF}(X, A) = \frac{\sum_{j=1}^n (\psi_X)_j (\psi_A)_j^*}{\sum_{j=1}^n (\psi_A)_j (\psi_A)_j^*}, \quad (1a)$$

and there are two possible expressions relating the two mode shapes, depending upon which one is taken as the reference,

$$\text{MSF}(A, X) = \frac{\sum_{j=1}^n (\psi_A)_j (\psi_X)_j^*}{\sum_{j=1}^n (\psi_X)_j (\psi_X)_j^*}. \quad (1b)$$

It should be noted that this parameter gives no indication as to the quality of the fit of the points to the straight line, simply its slope.

The second parameter is referred to as a Mode Shape Correlation Coefficient (MSCC) or more popularly the Modal Assurance Criterion (MAC) and provides a measure of the least-squares deviation or ‘scatter’ of the points from the straight line correlation. This parameter is defined by,

$$\text{MAC}(A, X) = \frac{\left| \sum_{j=1}^n (\psi_X)_j (\psi_A)_j^* \right|^2}{\left\{ \left(\sum_{j=1}^n (\psi_X)_j (\psi_X)_j^* \right) \cdot \left(\sum_{j=1}^n (\psi_A)_j (\psi_A)_j^* \right) \right\}}, \quad (2)$$

or

$$\text{MAC}(A, X) = |\{\psi_X\}^T \{\psi_A\}|^2 / [(\{\psi_X\}^T \{\psi_X\}) \cdot (\{\psi_A\}^T \{\psi_A\})],$$

and is clearly a scalar quantity, even if the mode shape data are complex. In the same way that the modal scale factor does not indicate the degree of correlation, neither does the Modal Assurance Criterion discriminate between random scatter being responsible for the deviations or systematic discrepancies, as described earlier. Thus, whereas these parameters are useful means of quantifying the comparison between two sets of mode shape data, they do not present the whole picture and should preferably be considered in conjunction with the plots of the form shown in figure 3.

It is worth considering two special cases: (i) where the two mode shapes are identical and (ii) where they differ by a simple scalar multiplier. Thus in case (i), we have,

$$\{\psi_X\} \equiv \{\psi_A\},$$

for which it can be seen that

$$\text{MSF}(X, A) = \text{MSF}(A, X) = 1,$$

and also that,

$$\text{MAC}(X, A) = 1.$$

In the second case, (ii), we have $\{\psi_X\} = \gamma\{\psi_A\}$ and we find that

$$\text{MSF}(X, A) = \gamma, \quad \text{while} \quad \text{MSF}(A, X) = 1/\gamma,$$

although, since the two modes are still perfectly correlated, we still have,

$$\text{MAC}(X, A) = 1.$$

In practice, typical data will be less ideal than this and what is expected is that if the experimental and theoretical mode shapes used are in fact from the same mode, then a value of the assurance criterion of close to 1.0 is expected, whereas if they actually relate to two different modes, then a value close to 0.0 should be obtained. Given a set of m_X experimental modes and a set of m_A predicted modes, we can compute a table of $m_X \times m_A$ Modal Assurance Criteria and present these in a matrix which should indicate clearly which experimental mode relates to which predicted one. Such a table is shown in figure 4a, together with some of the common graphical presentations used to display these data, in figures 4b and c. It is difficult to prescribe the precise values which the assurance criterion should take in order to guarantee good results. Generally, it is found that a value in excess of 0.9 should be attained for well-correlated modes and a value of less than 0.1 for uncorrelated modes. In some situations, the boundaries for ‘acceptable’ and non correlation are quoted as above 80% and less than 20% respectively. However, the significance of these quantities depends considerably on the specific data points used in the correlation (see below) and on the subsequent use planned for the model – some are much more demanding than others – and so considerable caution should be used in attaching quantitative significance to the absolute values of MAC obtained in practical cases. The greatest value of these coefficients lies in their use for comparison purposes.

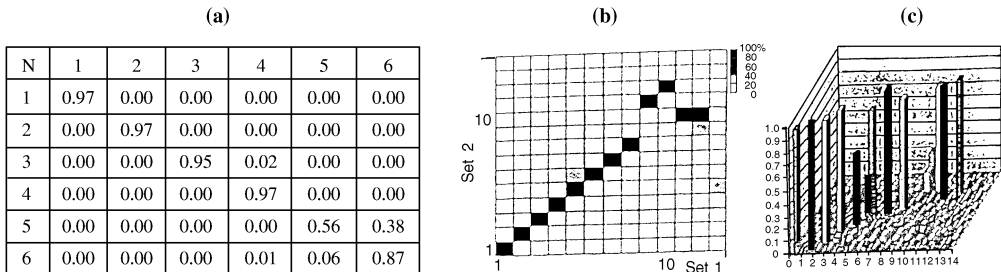


Figure 4. Alternative presentations of MAC.

It is worth noting some of the causes of less than perfect results from these calculations. Besides the obvious reason, that the model is incorrect, values of the MAC of less than unity can be caused by,

- (i) nonlinearities in the test structure,
- (ii) noise in the measured data,
- (iii) poor modal analysis of the measured data, and
- (iv) inappropriate choice of DOFs included in the correlation.

2.2d *Features of the MAC – AutoMAC*: In view of the widespread use of the MAC, and of the dangers of its misinterpretation, it is appropriate to include here a fuller discussion of the features of this approach.

The first consideration concerns the choice of DOFs which are included in the calculation. It is clear that if all the degrees of freedom in the model are included then a very accurate measure of the correlation between the two vectors will result. However, it is also clear that a different result will be obtained for the MAC if only a few of the full set of DOFs are included.

In practical terms, there is a difficulty in deciding how many, or – more accurately – which, DOFs need to be included in order to avoid the spatial aliasing problem. A detailed procedure for answering this question is given elsewhere but here it will suffice to show how the MAC can be used to check whether a given selection of DOFs is adequate or not. This is done via a version of the MAC called the AutoMAC in which a set of mode shape vectors are correlated *with themselves*. If, for example, we take the mode shape vectors for the analytical model but defined only at the DOFs which are to be used in the correlation with the experimental model (i.e. those DOFs which are included in the modal test) and compute the MAC table with themselves, we produce a result such as those shown in figures 5a and b. In the first of these, figure 5a, a ‘full’ set of DOFs is included while in figure 5b only the reduced subset that was used in the above example is included. From these AutoMAC plots we observe a number of features: (i) all the diagonal values are identically unity – they must be 100%, by definition because each mode shape must correlate perfectly with itself; (ii) the AutoMAC matrix is symmetric, and (iii) there are a number of non-zero off-diagonal terms, which means that some of the modes exhibit a degree of correlation with others, a result not immediately expected, since the modes are supposed to be ‘orthogonal’. However, there are two reasons why this orthogonality property does not translate to a perfectly diagonal AutoMAC matrix. The first is because the orthogonality property is only strictly applicable when the mass matrix is used and second, because it is only applicable when *all* the DOFs are included in the calculation. We shall return to the question of the mass matrix later, but for the moment our interest is focused on the matter of the number of DOFs included. In the limit, it can be seen that if we only define each mode shape by the amplitudes at just 2 DOFs, then most modes would look very similar and, indeed, if we did a formal correlation, would be found to be highly correlated with roughly half of all the modes included in the process. It is necessary to include sufficient DOFs to ensure the effective discrimination between the various modes. In fact, it is necessary to include in the selection of measured DOFs only those which are required to ensure that the eigenvector submatrix $[\Psi_1]_{n \times m}$, which is formed from,

$$\begin{bmatrix} [\Psi_1]_{n \times m} & [\Psi_3]_{n \times (N-m)} \\ [\Psi_2]_{(N-n) \times m} & [\Psi_4]_{(N-n) \times (N-m)} \end{bmatrix},$$

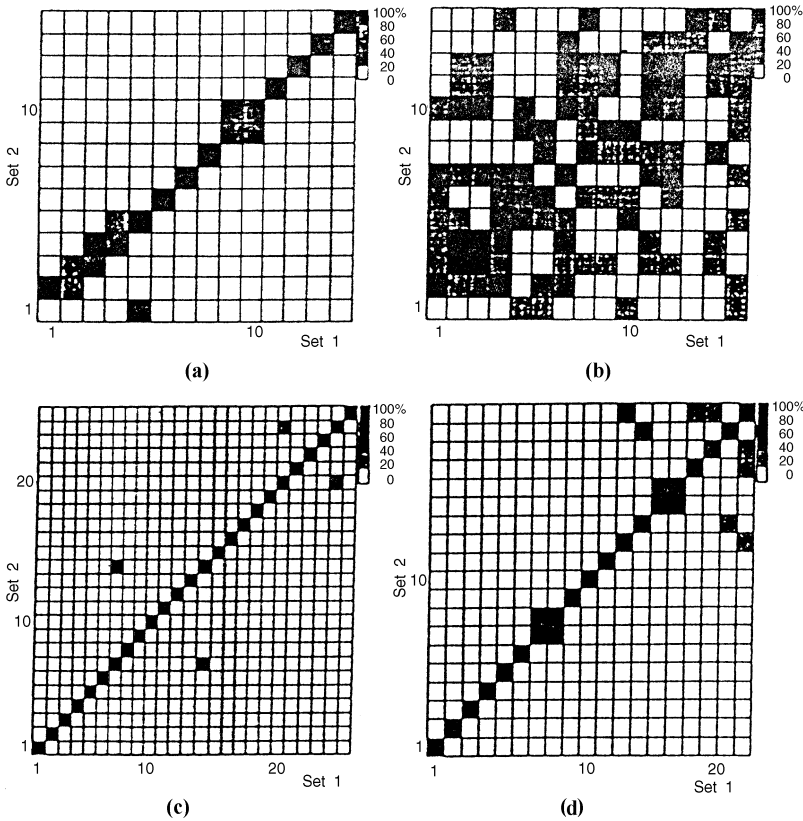


Figure 5. MAC tables – (a) acceptable result, (b) insufficient DOFs, (c), (d) AutoMAC to test adequacy of DOFs.

where n represents the included DOFs, and m the measured modes, is non-singular. Achieving this condition is not a trivial matter but the degree to which a given choice of DOFs satisfies it can be readily demonstrated by using the AutoMAC. Figures 5c and d show the suitability of the selected DOFs for an industrial structure, first on the FE model results and then on the experimental data themselves, both results confirming that the relatively small number of DOFs included in the correlation are suitable for the task of matching the correlated mode pairs.

Normalised MAC – Reference was made in the previous section to the absence of the mass matrix in the MAC calculation, an absence which means that the MAC is not a true orthogonality check. It is possible to remedy this limitation by including information which may be available on the mass of the system or, equally, of its stiffness, but to do so is relatively expensive, and constitutes a significant extension to the effort required to perform these checks. Bearing in mind the essentially comparative nature of the MAC coefficients, this extra effort is seldom warranted. However, in more advanced cases, including those where an automatic correlation procedure is sought, and where the numerical values of the correlation coefficients are likely to be used in subsequent stages of validation, the extension of the concept to the mass-normalised version may be considered. The formula for this version of the MAC, sometimes referred to as the Normalised Cross Orthogonality

(NCO), is given by,

$$\text{NCO}(A, X) = \frac{|\{\psi_X\}^T [W] \{\psi_A\}|^2}{[\{\psi_X\}^T [W] \{\psi_X\}] \cdot (\{\psi_A\}^T [W] \{\psi_A\})},$$

where the weighting matrix, $[W]$, can be provided either by the mass or stiffness matrices of the system.

The main difficulty to be overcome, even in those cases where a full mass and/or stiffness matrix is available from the analytical model, is that of reducing or condensing this mass matrix to the order of the specific DOFs for which data are available. A Guyan-type or equivalent reduction must be made if the mass matrix is to be used explicitly. One of the more practical approaches uses the SEREP-based reduction process. In this approach, a pseudo-mass matrix of the correct size is computed from the simple formula,

$$[M^R] = [\Psi]^+{}^T [\Psi]^+,$$

using either the limited measured eigenvectors or the corresponding analytical ones (preferred because of their greater accuracy). This pseudo-mass matrix can then be used in the NCO calculation as a weighting matrix and a readily-accessible version of a Normalised MAC – sometimes referred to as the SEREP-Cross-Orthogonality (SCO) coefficient – is thereby computed. Examples of both MAC and of SCO are shown in figure 6.

There are one or two other related issues which concern the question of choice of DOFs. It must be remembered that the accuracy of amplitude measurements made with attached accelerometers can be subject to considerable errors in cases where the motion in directions perpendicular to the axis of actual measurement is markedly greater than that being measured. In these circumstances, errors of 100% or more in the (small) amplitudes being recorded are common, and can easily contaminate correlation calculations as a result. Also, the difference between the units used in translational and rotational DOFs means that if both types of response are included in such a calculation, then one or other of these two sets of DOFs will be weighted quite differently to the other. Thirdly, there is the important matter of the exactness with which the two sets of DOFs match each other. The precise location of the measurement sites in relation to the FE model node points can also be a critical feature. It is shown in the next section just how relatively slight discrepancies in the location of a DOF can influence the numerical value of the recorded mode shape amplitude, and of the resulting correlation coefficients.

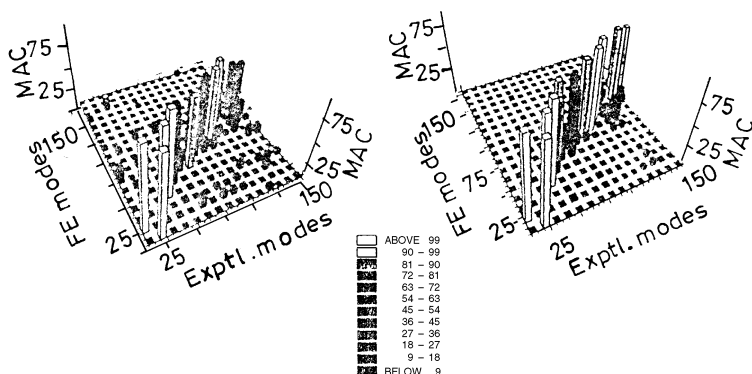


Figure 6. Plots of MAC and SCO.

These, and other, considerations show how sensitive the correlation calculations can be to the choice and accuracy of the mode shape amplitude data which are used in the process.

2.2e *COMAC*: The preceding paragraphs have all been concerned with the influence on the correlation process of the individual DOFs which are involved in the calculations. While these DOFs do not appear explicitly in any of the MAC coefficients, their importance is clearly seen by comparison of the values produced using different selections of DOFs. Clearly, there is a spatial dependence of the correlation parameters and our goal in the present section is to seek a way of expressing that dependence directly, so that a measure of the degree of correlation is presented as a function of the individual DOFs. This goal can be realised by rearranging the order in which the correlation calculations are performed and by defining a quantity called the ‘Coordinate MAC’ or COMAC.

In the calculation of the MAC between two vectors, a summation is made over all the DOFs included, resulting in a single coefficient for that pair of modes. The first step in the calculation of the COMAC is to preserve the individual elements in that summation, noting that each refers to one particular DOF. If we then take another pair of vectors, or modes, from the same two sets and repeat this step, we arrive at a second set of individual terms, relating to the same set of DOFs, and so on for as many mode pairs as we choose to include. If we restrict the pairs of modes thus included to the already-identified correlated mode pairs, then the data we have gathered in this way contain information about the quality of the correlation between properly matched mode shape vectors and so we can use it to define this correlation in more detail. In effect, the MAC value for each of the selected mode pairs is obtained by summing the contributions for all CMPs while a summation down each individual column yields information about the degree of correlation observed for that individual DOF. Suitably normalised to present a value between 0 and 1, the COMAC parameter for an individual DOF, i , is expressed as,

$$\text{COMAC}(A_i, X_i) = \sum_{l=1}^L |(\psi_X)_{il}(\psi_A)_{il}|^2 / \left(\sum_{l=1}^L (\psi_X)_{il}^2 \cdot \sum_{l=1}^L (\psi_A)_{il}^2 \right).$$

Here, l represents an individual correlated mode pair, of which a total of L are available, where L may well be less than the total number of modes in both sets, m_A for the analytical model or an m_X being the total number of experimentally-determined modes.

The COMAC can be displayed in different ways, the most obvious being simply a diagram of its value against the DOF number, as shown in figure 7. Alternatively, it is




Figure 7. Plots of COMAC values for all DOFs.

possible to use a display of the actual structure as the basis for a diagram, a format which gives a much more graphic illustration of the result.

As with many of these parameters, the correct interpretation of COMAC can be difficult to make. One is tempted to conclude that regions of the structure which show up as having (relatively) low values of COMAC are those regions which harbour the discrepancies which are responsible for the differences observed between the two models. This is seldom the case for the simple reason that regions of low COMAC correspond to regions where the *consequences* of any discrepancies between the two models are felt, rather than where they are actually located. Thus, large drops in COMAC are often observed at regions of large amplitude, such as at the free ends of beams, where the effects of inaccurate flexibility data in other parts of the structure are most dramatically felt. Care and ingenuity must be exercised in making such interpretations but the fact remains that the existence of systematic patterns of COMAC values almost always indicate systematic sources of discrepancy between the two models, even if these are not immediately located.

2.3 Comparison of response properties

2.3a Comparison of individual response function: If we start with the experimental model, we find that the raw data available in this case are those describing the time histories of the excitation and response properties of the test structure during the measurement. Although it is true that these data constitute the most direct measurement of the structure's actual dynamic behaviour, it is difficult to make comparisons between these properties and the corresponding quantities computed from the analytical model. This is so for several reasons but the most important of which is that the actual time histories are very sensitive to certain properties in the analytical model which are very difficult to estimate, most critically, the damping, but also a range of other features which, in themselves, are not critical but which combine to make useful comparisons of time histories difficult to achieve. As a result, attempts to extract useful information from comparisons between prediction and observation of these raw response data are not usually made.

The next level of proximity to the actual measurements (a condition which is important to achieve if an honest comparison of observation and prediction is to be made) is provided by the response functions which are derived from spectral analysis and further processing of the original response measurements. These response functions are generally presented as FRFs, or sometimes IRFs, and it is on these that we shall focus our efforts for further comparison procedures. In its simplest form, this level of comparison is made with an individual response function, and is shown by overlaying the measured curve on its analytically predicted counterpart, although it must be borne in mind at the outset that two important estimates will have to be made in order to be able to compute the theoretical curve, the nature and level of damping, which is not usually part of the initial modelling process, and the number of modes which will be included in the summation made to compute the response functions. The first of these two estimated parameters affects the predicted FRFs only in the immediate vicinity of resonances or antiresonances by limiting the sharpness of their peaks and troughs, but the second parameter can have more significant effects on the general shape of the curves in all regions away from the resonances if sufficient modes are not included. Strictly, a check that a sufficient number of modes have been included should be made before any response function comparisons are attempted.




Figure 8. Comparison of individual FRFs. (measured (point 37) ...; predicted (point 2) —, (point 3) - -; forcing point 11.)

A second example is shown in figure 8 where a transfer mobility for a different structure is illustrated, again for both experimental and predicted data. However, in this case the location of the response point used in the modal test does not coincide exactly with any of the mesh of grid points used in the analytical model, thereby making a direct comparison impossible. In order to proceed, the predicted curves relating the two grid points closest to the test position (and these were only a few mm away, on a plate-like structure of some $1\text{m} \times 2\text{m}$) are used and are shown in figure 8. In this example, it is clear that not only are there marked differences between the two models (albeit of a different type to the previous case), but also there are striking differences between the two predicted curves which relate to two points very close to each other on the structure. This last observation is very important when we consider how to assess the degree of correlation between the experimental and predicted models. Because the particular parameter being measured (an FRF) can be very sensitive to the exact location of the response point (and, possibly, to the excitation point, although that does not suffer from the same difficulty as does the response in the example cited), major differences may be apparent at the comparison stage which do not directly reflect on the quality of the model, but on something much more basic – namely, the coordinate geometry used in both instances.

2.3b Comparison of complete sets of FRFs: It can be noted that in a typical modal test a set of FRFs is measured consisting of at least one column (i.e one vector) in the FRF matrix based on the measured DOFs, and sometimes including data from several such columns (vectors). It is thus possible to envisage the curve/curve comparison described above having to be applied to a large number of such data in order to gain an overall impression of the degree of correlation between measurement and prediction. This is a daunting task, and difficult to perform effectively because there is simply too much information for the analyst to retain and to sort. It is not unlike the problem faced by the need to compare several mode shapes simultaneously, except that in this case, there are of the order of 400 or 800 such vectors because there is one for each excitation frequency used in defining the FRFs. This observation leads to the idea of applying the MAC approach to the correlation of two vectors, one from the measured data and the other from a corresponding analytical model prediction. Thus we can define a frequency response assurance criterion, or FRAC, as follows.

$$\text{FRAC}(A(\omega_j), X(\omega_i)) = \frac{|\{H_X(\omega_i)\}^T \{H_A(\omega_j)\}|^2}{(\{H_X(\omega_i)\}^T \{H_X(\omega_i)\}) \cdot (\{H_A(\omega_j)\}^T \{H_A(\omega_j)\})}$$

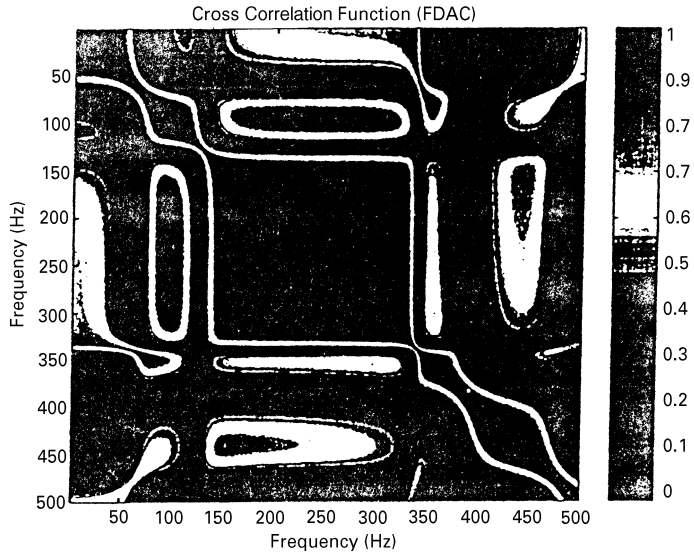


Figure 9. Correlation of individual FRFs (FRAC).

Clearly, a diagram of the type previously used for the MAC can be used to display this function, although it will have a much denser form as a result of the large number of frequencies (typically, several hundreds) by comparison with the usual number of modes (typically, tens) as illustrated in the example shown in figure 9.

3. Concluding remarks

In this paper, we have reviewed a wide range of methods for the comparison and quantitative correlation between the dynamic properties predicted by a theoretical model and those measured in a modal test. At the conclusion of such correlation procedures, it is possible to determine correlated mode pairs (a prerequisite of most updating methods) and to provide some early indications of the location of errors.