Review

Emerging Trends in Optimal Structural Health Monitoring System Design: From Sensor Placement to System Evaluation

Robert James Barthorpe * and Keith Worden

Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, Mappin Street, Sheffield S1 3JD, UK; k.worden@sheffield.ac.uk

* Correspondence: r.j.barthorpe@sheffield.ac.uk

Received: 31 March 2020; Accepted: 24 June 2020; Published: 1 July 2020

Abstract: This paper presents a review of advances in the field of Sensor Placement Optimisation (SPO) strategies for Structural Health Monitoring (SHM). This task has received a great deal of attention in the research literature, from initial foundations in the control engineering literature to adoption in a modal or system identification context in the structural dynamics community. Recent years have seen an increasing focus on methods that are specific to damage identification, with the maximisation of correct classification outcomes being prioritised. The objectives of this article are to present the SPO for SHM problem, to provide an overview of the current state of the art in this area, and to identify promising emergent trends within the literature. The key conclusions drawn are that there remains a great deal of scope for research in a number of key areas, including the development of methods that promote robustness to modelling uncertainty, benign effects within measured data, and failures within the sensor network. There also remains a paucity of studies that demonstrate practical, experimental evaluation of developed SHM system designs. Finally, it is argued that the pursuit of novel or highly efficient optimisation methods may be considered to be of secondary importance in an SPO context, given that the optimisation effort is expended at the design stage.

Keywords: sensor placement optimisation; structural health monitoring; damage identification; mutual information; evolutionary optimisation

1. Introduction

The broad aim of Structural Health Monitoring (SHM) is to be able to objectively quantify the condition of an engineered structure on the basis of continually or periodically observed data, enabling incipient faults that may ultimately lead to failure of component or system to be detected at an early stage. These monitoring outcomes may be subsequently used to inform management of the structure, including decisions on how and when to deploy corrective actions and maintenance. Achieving high quality outcomes from the SHM process critically depends upon the quality of information gathered from the measurement system deployed upon the structure, which typically takes the form of a network of acceleration and/or strain transducers and appropriate acquisition hardware. Thus, the question of where to place the available sensors is key to achieving these desired outcomes. In most scenarios it is impractical to instrument every location of interest within a given structure. This is of particular concern in applications where the costs that are associated with deploying and maintaining a sensor network are high, for example in the aerospace sector where the impact of any mass addition on fuel efficiency must be considered. The desire to maximise the value of information returned by a reduced sensor set has led to the development of formal methods for sensor placement optimisation

(SPO) for damage identification. As the size and complexity of the structures considered for monitoring has increased—with typical target applications ranging from aircraft to offshore wind turbines and large civil structures—so has the desire for effective methods for evaluating and optimising the design of SHM systems. The advent of cost effective and reliable wireless sensor networks adds a further consideration on this front.

This paper acts as an update on an earlier manuscript [1] and seeks to present the aims of SPO in an SHM context, set out the commonly adopted frameworks for tackling this problem and provide a critical review of the literature in this area. Most importantly, the particular contribution of this paper is to identify the key emerging topics that represent directions for future study. Given the wealth of literature in this area, it would be unwise to claim that this review is truly comprehensive; the quantity of (often closely related) works is simply too great. Instead, the paper is structured in order to highlight key themes within the literature and, consequently, demonstrate interrelations between approaches and avenues to novelty for researchers in this area.

A number of review articles exist on the topic of SPO in structural dynamics. Yi and Li [2] presents an overview of cost functions and optimisation methods focusing on civil infrastructure applications. A recent review article by Ostachowicz et al. [3] provides an update on the literature in this area with emphasis placed on methods appropriate to guided-wave SHM. A particular focus of the review in ref. [3] was to describe the numerous metaheuristic methods that have been applied in an SPO context; these are critiqued in Section 4.2. While a degree of overlap between reviews of this type is inevitable, a key distinguishing element of the current contribution is the emphasis on the identification of emerging methods that may allow the design of robustly-optimised SHM systems for which reliable a priori measures of expected benefit may be determined.

The layout of the paper is as follows. Section 2 provides an overview of the SPO problem as it applies to SHM. Section 3 introduces the cost functions. In Section 4, the optimisation methods that have been applied to the SPO task are summarised. Emerging trends are identified in Section 5 prior to concluding comments that are presented in Section 6.

2. Overview of the Sensor Placement Optimisation Problem

In a structural dynamics context, sensor placement optimisation first emerged for tasks such as the modal analysis of structures. Prior to the adoption of formal techniques, achieving successful modal analysis outcomes was highly reliant upon the experience and insight of the dynamicist conducting the test. By making use of practical steps, such as prioritising expected antinode location and avoiding expected nodes, it would be possible to create acceptable ad hoc sensor distributions that largely served well for distinguishing and identifying modes of non-complex structures. In cases where resource allowed, it would have been possible to trial different sensor configurations experimentally prior to adopting a final layout. However, as the complexity of the structures being considered has grown, it has become increasingly common practice to formalise the sensor placement problem as one of optimisation, with a cost function derived and subsequently minimised while using an appropriately chosen optimisation technique. The data that are used to conduct this optimisation may be drawn directly from experiment or, more commonly, may be based on the predictions of a numerical model. In parallel to this move to more formal optimisation methods has been the growth of research into SHM systems. While early SHM methods were largely based upon modal data (for example, changes in natural frequencies and modeshapes), recent years have seen a growth of methods making use of a far more diverse set of features, broadening the information sought from the SHM sensor network.

While the SPO problem may, in principal, be posed as one with a continuous outcome (e.g., returning the cartesian coordinates of individual sensor locations within a continuous space), the approach adopted almost universally in the reviewed literature is to pose the SPO task as one of combinatorial optimisation. This takes as it's starting point a candidate set of size n that contains all degrees of freedom (DOFs) that are available as sensor locations. The SPO task is to reduce this set to a smaller measurement set of size m that comprises only those DOFs to be employed as sensor

locations. The number of combinations available for a given measurement sensor set size m is given by the binomial coefficient,

$$\binom{n}{m} = \frac{n!}{m!(n-m)!} \quad (1)$$

Searching for globally-optimal solutions becomes computationally demanding for anything other than cases with small n . Figure 1 provides an illustrative example of an SPO application for SHM, and it serves to demonstrate this point. The objective of the SHM system in the presented example is to detect damage introduced into a composite aircraft wing structure via changes in the acceleration response spectra recorded at a set of measurement locations. For any given measurement configuration, a set of damage-sensitive features may be derived from the retained spectra. Representations of this feature set from both the damaged and undamaged state of the structure may subsequently be used to train a damage detector. In this example, a candidate set of 36 piezoelectric accelerometers are mounted on the upper surface of the wing, with the system excited using an electrodynamic shaker mounted towards the wingtip. Thus, the SPO task is to select the measurement sensor set that generates the most robustly discriminative damage detector.

Figure 1. A composite wing structure mounted in a free-free configuration. A candidate set of 36 accelerometers are mounted to its upper surface.

In the experimental example that is presented in Figure 1, the candidate set comprises $n = 36$ sensor locations. Setting the measurement set size to $m = 6$ sensors would result in 1.95×10^6 possible combinations to consider. Thus, arriving at an optimal measurement set in this scenario may become a computationally expensive task, and particularly so if the evaluation of a given measurement set involves tasks, such as training and testing a classifier. This issue grows in importance as n increases. To illustrate this, consider the common case that SPO for the structure shown in Figure 1 is conducted while using the predictions of finite element analysis (FEA) rather than data from the physical structure. If, for the purposes of illustration, $n = 1000$ nodal points were adopted as admissible candidate locations, the number of measurement set combinations for the same $m = 6$ scenario would grow to 1.37×10^{15} . This fact prompts interest in methods for efficient combinatorial optimisation, with an overview of applied techniques being given in Section 4.

3. Cost Functions

This section seeks to provide an overview of the key approaches that have been proposed for sensor placement in the structural dynamics context, building from foundational techniques, such as the effective independence (EI) and kinetic energy (KE) methods, towards more recent methods based

upon classification outcomes that are more specifically tailored to the SHM task. The intention of this section is to highlight both the development arc of these methods and the key considerations that are to be made by the SHM practitioner when seeking to select, develop, and implement them.

3.1. Information Theory

Several of the cost functions introduced in this section are based upon information theory; this is natural, given that the essence of the SPO problem is to extract the most useful information possible given a restricted number of available channels (or, equivalently, measured DOFs in the SPO case). Prior to introducing the cost functions that have been applied for SHM systems, there is value in briefly reviewing relevant concepts.

3.1.1. Fisher Information

Fisher information offers a measure of the information that a sampled random variable X contains about an unknown parameter θ . Formally, Fisher information is based upon the score i.e. the gradient of the log-likelihood function, $\ln f(X; \theta)$, relating X , and θ . This gradient is usually expressed as the partial derivative of this function w.r.t. to θ . The Fisher information itself is then the variance of the score,

$$I(\theta) = E \left[\left(\frac{\partial}{\partial \theta} \ln f(X; \theta) \right) \middle| \theta \right] \quad (2)$$

where $E[\cdot]$ denotes the conditional expectation, in this case with respect to θ . In the case that there are multiple unknown parameters $\theta = [\theta_1 \dots \theta_N]$, the Fisher information may be stated in matrix form with elements,

$$[I(\theta)]_{ij} = E \left[\left(\frac{\partial}{\partial \theta_i} \ln f(X; \theta) \right) \left(\frac{\partial}{\partial \theta_j} \ln f(X; \theta) \right) \middle| \theta \right] \quad (3)$$

Techniques that are based upon maximising the determinant of the Fisher information matrix (FIM) have been widely used as a basis for sensor placement optimisation. In these approaches, the available data, X , is typically taken to comprise the measurements available from a given sensor distribution. Various quantities may then be adopted for the unknown, sought parameters θ . For example, in the case of Effective Independence (described in Section 3.2.1) the quantities that are chosen for θ are the elements of the modeshape matrix.

3.1.2. Mutual Information

Mutual information gives a measure of the mutual dependence between two random variables. The mutual information between two variables may be expressed as,

$$I(x, y) = \log_2 \left[\frac{p_{X,Y}(x, y)}{p_X(x)p_Y(y)} \right] \quad (4)$$

where x and y are realisations of the jointly-distributed random variables X and Y , $p_{X,Y}$ is the joint probability density function of X and Y and p_X and p_Y are the marginal densities of X and Y .

For the SPO problem, mutual information may be interpreted as how much information one can “learn” about a sensor location from any another. If there are two sets of measurement locations, $\mathbf{A}$ and $\mathbf{B}$, the amount of information that is learnt by sensor location a_i about location b_j is represented by the mutual information $I(a_i, b_j)$. In contrast to many applications of mutual information in machine learning, the approach that is taken in sensor placement optimisation is typically to minimise the mutual information between sensor locations, thus promoting independence. In the case that a_i and b_j are completely independent of one another, $I(a_i, b_j)$ drops to zero. A typical approach making use of sequential sensor placement (see Section 4.1) would involve $\mathbf{A}$ comprising a set of adopted measurement locations and $\mathbf{B}$ comprising a set of locations being considered for addition to it. The average mutual information between the adopted set of locations in $\mathbf{A}$ and each candidate for

addition in **B** may be calculated by evaluating I for in a pairwise fashion and taking an average. The candidate location from set **B** that is determined to possess the lowest average mutual information with those in **A** is then adopted at each step.

3.1.3. Information Entropy

Information entropy offers a measure of the uncertainty in a model's parameter estimates. Optimal sensor placement might be achieved by minimising the change in the information entropy $H(D)$ given by,

$$H(X) = -E [\ln \mathcal{L}(X|\theta)] \quad (5)$$

where θ is the uncertain parameter set, X represents the experimental test data, and $E(\cdot)$ denotes the expectation with respect to θ . A rigorous mathematical description is given in ref. [4], where it is shown that, for large datasets, the information entropy depends on the determinant of the FIM.

3.2. Modal Identification Based Cost Functions

The aim of the sensor placement task in a modal identification context is to select a measurement location set from a large, finite candidate set, such that the modal behaviour of the system remains as accurately represented as possible. Successful modal identification from a reduced sensor set requires three key decisions to be made, as discussed by Udwadia [5]:

1. Sensor quantity: how many sensors are required to enable successful modal identification?
2. Sensor placement: where should the available sensors be placed in order to best capture the required data?
3. Evaluation: how may the performance of the final, optimised sensor configuration be quantified?

Bounds on sensor quantity will generally be apparent from early in the test planning process. The lower bound is set by the requirement that for the modeshapes to be uniquely identifiable, the sensor quantity must be at least equal to the number of modes of interest. The upper bound will usually arise from resource limitations e.g., costs per channel or availability of equipment. Between these bounds, there may be some flexibility to consider the addition of sensors beyond the lower bound in order to allow for either greater ability to visualise modeshapes, or to provide a degree of robustness to sensor failure. Sensor placement and system evaluation are closely related and constitute the primary focus of this article. For the sensor placement step, the key elements are to adopt an appropriate performance measure/cost function (the focus of the present section); and then to select an appropriate method with which to optimise it (covered in Section 4). Evaluation of the performance of the optimal measurement set is an important, final gating activity prior to moving on to testing, enabling the expected performance of the proposed sensing system to be quantitatively evaluated and compared against test requirements. A distinction is drawn here between the evaluation methods used as part of the sensor placement optimisation step and those adopted for the final system evaluation step, as these may not necessarily be the same. It may, for example, be the case that SPO is conducted using a computationally inexpensive abstraction of the procedure applied at the final system evaluation step.

Note that the three decisions that are listed above will not necessarily be made in sequential steps; in general, there will be interplay and iteration between them. Also note that, while initially proposed in a modal identification context, these decisions are of no less importance when considering SPO for SHM.

3.2.1. Effective Independence (EI)

The effective independence (EI) method was introduced by Kammer [6] following earlier work by Shah and Udwadia [7]. The method makes use of the Fisher Information Matrix, with the modeshape

matrix being adopted as the quantity of interest. Maintaining the determinant of the FIM leads to the selection of a set of sensor locations for which the modeshapes are as linearly independent as possible.

Central to the method is the EI distribution vector E_D , which calculates the contribution that is made by each individual sensor location to the rank of the prediction matrix E . E can only be full rank if the mode partitions resulting from a given measurement location set are linearly independent. The distribution vector E_D is defined as the diagonal of E ,

$$E = [\Phi] \{ [\Phi]^T [\Phi] \}^{-1} [\Phi]^T \quad (6)$$

where Φ represents the mass normalised modeshape of the system realised at all retained response locations. The algorithm is sequential in nature. At each step, the terms in E_D are ranked according to their contribution to the determinant of the FIM. The lowest ranked sensor location is identified and deleted from the candidate set, as are the corresponding elements in the modeshape matrix. The new, reduced sensor set is then re-ranked, and the process is repeated in a sequential manner until the desired number of sensors is obtained.

The work was extended in ref. [8] to consider the effect of model error on EI outcomes, and in ref. [9] to consider the effect of measurement noise. The application of a Genetic Algorithm (GA) based approach is presented in ref. [10]. An alternative approach to the EI technique is presented in ref. [11]. Instead of sequentially removing sensors from a candidate set, the sensor set is sequentially expanded until the required number of sensors is achieved. The method selects the sensor location that offers the greatest increase in the determinant of the FIM. This expansion can begin with a single sensor, and multiple locations can be added at each step of the iteration if required. This change in approach reduces the issue of computational expense associated with large candidate sets, and presents the possibility of specifying a desired (or existing) sensor configuration at the outset, to which further sensors may be optimally added. The approach was demonstrated for 27 target modes provided by a large FEA model of an aerospace vehicle. The results from the new expansion method and from a conventional EI reduction method were compared. From a candidate set of 29,772 locations, 389 were chosen as measurement locations, with agreement between the two approaches for all but 9 sensors. A significant saving in computational costs was made. A commonly observed drawback of the EI algorithm is that it does not penalise sensor locations that display low signal strength, thus making it susceptible to poor performance in noisy conditions.

Penny et al. [12] compared the performance of the EI algorithm against that of a scheme based upon classic Guyan reduction using an *a priori* FE model of a cantilever beam as a case study. The Guyan reduction approach proceeds by sequentially removing measurement locations at which the inertial forces are small in comparison to the elastic forces. The original motivation for Guyan reduction was to allow for large finite element models to be reduced to a smaller set of master DOFs prior to eigenvalue analysis, with the authors postulating that the behaviour that made particular DOFs suitable for adoption as master nodes would also make them good candidates for inclusion in an experimental measurement set. The measurement sets that were proposed by each method were compared using three evaluation measures: the size of the off-diagonal elements of the modal assurance criterion (MAC) matrix; the condition number of a singular value decomposition (SVD) of the modeshape matrix; and, the determinant of the FIM. While the off-diagonal elements of the AutoMAC matrix should drop to zero in the case that the modes are not correlated, the SVD measure offers a more explicit indication of the linear independence of the modeshape vectors. It was found that the EI method performed better in the case that rigid body modes were present, while the Guyan reduction method worked well for cases where the structure was grounded and, thus, displayed no rigid body modes. However, the key conclusion was that, while both methods produced acceptable results, neither reached a global optimum due to the sequential deletion method adopted for optimisation. Sequential sensor placement schemes—and alternatives—will be returned to in Section 4.

3.2.2. Modal Kinetic Energy (MKE)

The MKE method [13] seeks to prioritise sensors that are located at points of maximum kinetic energy for the modes of interest, on the assumption that these locations will maximise the observability of those modes. MKE indices are calculated for all of the candidate sensor locations i , as follows

$$MKE(i) = \sum_r \left[\phi_{ir} \sum_j M_{ij} \phi_{jr} \right] \quad (7)$$

where ϕ is the target modeshape matrix, M is the mass matrix, r refers to the r th mode, and i, j to the i th and j th DOF, respectively. Thus, the modal displacements that are associated with each location are weighted by the corresponding component from the mass matrix M , providing a measure of the kinetic energy contribution of each location to the modes of interest.

The major advantage of the MKE method is that it promotes the placement of sensors in a location of high signal strength and so is less susceptible than EI methods to issues arising from low signal-to-noise ratios. However, this comes at the cost of not explicitly considering the independence of the returned modeshapes. Also note that the use of the mass matrix effectively limits the use of MKE to model-based approaches; down-selection from an experimental set is not an available option in this case.

Li et al. [14] explored the inherent mathematical connection between the EI and KE methods. It is shown that the first iteration of the EI method will always give the same result as the KE method, and that for the special case of a structure with an equivalent identity mass matrix the EI approach is an iterated version of the KE approach, with the reduced modeshapes re-orthonormalised at each iteration of the EI method. For cases with non-identity equivalent mass matrix the KE method incorporates the mass distribution that is associated with each candidate DOF as a weight, whereas the EI method is not mass dependent. Both methods are applied to data from the I-40 bridge in Albuquerque, New Mexico for three cases. From a candidate set of 26 locations, 25 locations are selected in Case 1, six are selected in Case 2, and eight are selected in case 3. For Case 1, the same sensor is selected for deletion by both methods, and the ranking sequence is identical. For Case 2, the same six sensors are selected and ranked in the same order by the two methods, but with different ranking values. In Case 3, there is agreement for seven of the eight required sensors, but discrepancies in the ranking values and ranking order.

3.2.3. Average Driving Point Residue (ADPR)

An alternative measure of the contribution of a given sensor location to the overall modal response of a system is given by the average driving point residue (ADPR) [15]. The intention of the ADPR is to provide a measure of the average response of the system at a given location to a broadband input that excites all modes of interest. The ADPR is calculated from data or FE predictions, as

$$ADPR(i) = \sum_r \frac{\phi_{ir}^2}{\omega_r} \quad (8)$$

where i is the candidate sensor location under consideration, $r = 1, \dots, N$ are the modes of interest, ϕ_{ir} is the i th element of the r th modeshape, and ω_r is the r th modal frequency. The ADPR in the form stated above provides a measure of average modal velocity, although minor variations allow modal displacements or accelerations to be considered instead [16]. In common with the closely-related MKE method, an advantage of using the ADPR approach over EI methods is that it tends to select sensors in areas of high signal strength, although, once again, at the cost of sacrificing explicit consideration of modeshape independence. The key differentiator of the ADPR and MKE approaches is that the ADPR does not require knowledge of the mass matrix, which enables it to be more easily applied in purely experimental scenarios where a physics-based model is not available.

A minor adaptation of the EI approach was introduced by Imamovic [16]. In the Effective Independence Driving Point Residue (EI-DPR) approach, the EI and ADPR values for each candidate location i are combined to produce an EIDPR value

$$E_D^{EIDPR}(i) = E_D^{EI}(i)ADPR(i) \quad (9)$$

Combining the metrics in this way was found to promote the placement of sensors in regions of higher signal to noise ratio than for the EI method alone, while also resulting in comparatively uniform sensor configurations. The Eigenvector Product (EVP) is a further closely-related method [17]. As the name suggests, this metric involves simply taking the product of the eigenvector elements for a given candidate location across all modes of interest,

$$EVP(i) = \prod_r |\phi_{ir}| \quad (10)$$

where i is the candidate sensor location and ϕ_{ir} is the i th element of the r th modeshape. A maximum for this product is deemed to be a candidate measurement location.

3.2.4. Modeshape Sensitivity

Shi et al. [18] present one of the earliest sensor placement approaches specific to structural health monitoring. A structural damage localisation approach that is based on eigenvector sensitivity is adopted, and a sensor placement optimisation approach is developed that uses the same method. The FIM approach is applied to the sensitivity matrix that is to be used for damage localisation, with those degrees of freedom that provide the greatest amount of information for localisation retained. The prediction matrix in this case is given by

$$E = [F(K)] \{ [F(K)]^T [F(K)] \}^{-1} [F(K)]^T \quad (11)$$

where $F(K)$ represents a matrix of sensitivity coefficients of modeshape changes with respect to a damage vector comprising stiffness matrix changes. The sensor placement and damage localisation algorithms are demonstrated for numerical and experimental data, and they are shown to be effective in indicating probable locations of both single and multiple damage locations. For the experimental survey of an eight-bay, three-dimensional truss structure, a set of 20 optimised measurement locations is selected from a candidate set of 138 DOFs. Correlation between the analytical and test modeshapes gives MAC values of over 0.99 for the first five modes. These five modes are used for the successful localisation of damage represented by the unscrewing of truss elements.

Guo et al. [19] highlight the difficulties that may occur in solving the matrix proposed by Shi et al. [18]. The problem is reformulated as an objective function that is based on damage detection, to be solved while using a genetic algorithm. An improved genetic algorithm (IGA) is introduced to deal with some of the limitations of the simple GA. The results from the new algorithm are compared with those from the penalty function method and the forced mutation method for a numerical simulation. The numerical example studied is based upon the same six-bay, two-dimensional truss used in Shi [18]. The improved genetic algorithm is shown to converge in far fewer generations than the penalty function and forced mutation methods. Improved convergence characteristics and low sensitivity to the initial population used are also demonstrated.

3.2.5. Strain Energy Distribution

Hemez and Farhat [20] present a sensor placement study that is specific to damage detection based on strain energy distributions. The EI algorithm that is presented in ref. [6] is modified to allow the placement of sensors according to strain energy distributions. First, a strain energy matrix E is constructed,

$$E = [\Psi^T \Psi] \quad (12)$$

where $\Psi = C^T$ and C represents a Cholesky decomposition of the stiffness matrix K i.e., $K = C^T C$. The elements of the distribution vector E_D that lies on the diagonal of E thus represents the contribution of each sensor location to the total strain energy recorded by the system. The proposed method proceeds in the same manner as the original EI algorithm, with the sequential deletion of locations from the candidate set until the desired number of sensors is reached. The performance of the EVP and EI methods is compared in ref. [20] for a damage detection case study on an eight-bay truss structure. The objective for the sensor network is to facilitate optimal model updating in the presence of damage. It was found that the damage detection performance of the updated model was sensitive to the chosen sensor distribution. While both of the algorithms returned measurement sets that are capable of detecting damage, the update that is based on the EI measurement set contained some inaccuracies when reporting the damage locations.

3.2.6. Mutual Information

In ref. [21], the selection of optimal sensor locations for impact detection in composite plates is approached while using a GA to optimise a fitness function based upon mutual information. The mutual information concept is used to eliminate redundancies in information between selected sensors, and rank them on their remaining information content. The method is experimentally demonstrated for a composite plate stiffened with four aluminium channels, with an optimal set of six sensors selected from a candidate set of 17. The selected locations are found to lie close to the stiffening components. It is concluded that this result is consistent with the stiffened regions being the most challenging for the assessment of impact amplitude; the deflections under impact in these regions are much smaller than those observed in the unconstrained centre of the plate.

Trendafilova et al. [22] use the concept of average mutual information to find the optimal distance between measurement points for the sensor network. An equally-spaced configuration is assumed, and the sensor spacing varied in order to minimise the average mutual information between measurement locations. The approach is tested for a numerical simulation of a rectangular plate using a previously-established damage detection and location strategy, while using damage probabilities as the assessment criteria. The approach is tested for three spacing values: an initial set containing 36 sensors, the optimal set given by the mutual information method and containing eight sensors, and a third set containing six sensors where the spacing between locations is greater than the calculated optimal value. The results show that the initial and optimally-spaced sets are equally successful in locating damage. The set with a greater spacing between sensors performs significantly less well.

3.2.7. Information Entropy

Papadimitriou et al. [23] present a statistical method for optimally placing sensors for the purposes of updating structural models that can subsequently be used for damage detection and localisation. The optimisation is performed by minimising information entropy, a unique measure of the uncertainty in the model parameters. The uncertainties are calculated while using Bayesian techniques and the minimisation realised using a GA. The method is tested using simplified numerical models of a building and a truss structure. The results show that the optimal sensor configuration depends on the parameterisation scheme adopted, the number of modes employed, and the type and location of excitation.

3.3. SHM Classification Outcomes

The dominant paradigm in SHM in recent years is that based upon statistical pattern recognition. This involves training statistical classification and/or regression algorithms in either a supervised or unsupervised manner. Thus, it is natural to consider classification outcomes as the basis for constructing cost functions. Formally, this requires the selection of metrics on classifier performance to be developed and evaluated using either the same data used for training (referred to as recall) or

performance on an independent testing set [24]. Studies that adopt cost functions that are based on classification outcomes are reviewed in Section 4, alongside the optimisation methods applied.

For the simplest case of binary classification, metrics may be drawn from a confusion matrix of the form that is shown in Table 1.

Table 1. Exemplar confusion matrix for binary classification.

		Predicted Class	
		Positive	Negative
Actual Class	Positive	TP	FP
	Negative	FN	TN

In the confusion matrix, the quantities TP, FP, TN, and FN refer, respectively, to the number of True Positive, False Positive, True Negative, and False Negative classification outcomes that are observed for observations from a presented dataset. The principle quantities of interest that may be extracted from them are the probability of detection P_D (or true positive rate (TPR)), which is given by,

$$P_D = \text{TPR} = \frac{\text{TP}}{\text{TP} + \text{FN}} \quad (13)$$

and the probability of false alarm rate P_{FA} (or false positive rate (FPR)), which is given by,

$$P_{FA} = \text{FPR} = \frac{\text{FP}}{\text{FP} + \text{TN}} \quad (14)$$

A perfect classifier would return a P_D of 1 and a P_{FA} of 0.

In general terms, classifiers operate by computing a score that is related to class membership and then comparing this score to some predetermined classification threshold. Thus, the particular choice of threshold value is critical in controlling the trade-off between the true positive and false positive rates. The Receiver Operating Characteristic (ROC) is a useful tool for providing a comparison between the TPR and FPR for any given threshold level. By evaluating the ROC value, a discrete set of possible threshold values, data may be gathered with which to plot an ROC curve. This curve may be employed to (a) evaluate the discriminative ability of the data and (b) select a level for the threshold that is appropriate to the problem, and it is for the first of these functions that it is typically applied in an SHM SPO context.

An example of an ROC curve for a binary classifier is given in Figure 2. In this example, a “Negative” class label indicates that the structure is undamaged and a “Positive” class label indicates that damage has occurred. A classifier offering no discrimination would correspond to the line of no discrimination indicated on the diagonal. The further that the ROC curve for a given classifier lies above this line, the greater its discriminatory performance.

The area under the ROC curve (AUC) offers a further, quantitative method for summarising the level of discrimination offered by a classifier. The AUC returns a value in the range [0,1], where a value of 1 would indicate perfect classification performance. A classifier that assigned labels at random would return a value of 0.5, with its performance being as indicated by the line of no discrimination in Figure 2. The AUC might be interpreted as the probability that the considered classifier will rank a randomly selected positive instance higher than a randomly chosen negative instance [25].

Classification performance-based metrics may be extended to the localisation task, with an exemplar confusion matrix for a three class case presented in Table 2.

Figure 2. An example of a Receiver Operating Characteristic (ROC) curve for a moderately well-separated dataset. For comparison, the line of no discrimination is indicated by the dashed line along the diagonal.

Table 2. Exemplar confusion matrix for damage localisation.

		Predicted Location		
		A	B	C
Actual Location	A	TP _A	E _{BA}	E _{CA}
	B	E _{AB}	TP _B	E _{CB}
	C	E _{AC}	E _{BC}	TP _C

For damage localisation, the class labels (A, B, and C in the example) will typically refer to a discrete set of potential damage locations, with TP_A representing the number of true positive outcomes for class A, E_{BA} representing mislabelling of an observation from class A as being from class B and so on. The metric of most interest is the probability of correct classification (P_{corr}) over all K classes i.e., the sum of the diagonal quantities TP_k (with $k = 1...K$) divided by the total number of observations in the test set, which may be expressed as

$$P_{corr} = \frac{\sum_k TP_k}{\sum_k TP_k + \sum_{i \neq j} E_{ij}} \quad (15)$$

An alternative approach for defining metrics on classifier performance is to explicitly evaluate the performance of the proposed subset relative to that of the full set. Such a scheme is proposed in ref. [26] for a Neural Network classifier, where the normalised mean squared error (NMSE) between a desired (i.e., best possible) classifier response and that achieved by a classifier trained on a subset of measurement locations is used as the basis for a range of cost (alternatively “fitness”) functions. The NMSE is given by,

$$NMSE(\hat{y}_i) = \frac{100}{N_T \sigma_i^2} \sum_{j=1}^{N_T} (y_i(j) - \hat{y}_i(j))^2 \quad (16)$$

where i represents the i th output neuron, N_T is the number of training sets indexed by j , and σ_i^2 is the variance of the output y_i . In ref. [26], it was proposed that the performance of the SHM system (a classifier trained using a subset of measurement points) may be judged on either the average value

of the NMSE over the whole set of classifier outputs; or alternatively using the maximum of the set of output NMSEs.

Finally, cost functions that are specific to particular classifiers may be adopted. These include the margin of separation for Support Vector Machine (SVM) classification, for example.

3.4. Summary

The methods highlighted in this section follow a progression from those developed for the general sensor placement case (e.g., for the purposes of modal analysis), through to those that are specific to SHM. Two dominant groups of approaches are apparent: methods that are based upon variations on the Fisher information matrix (and other information metrics); and, methods based upon maximising classification outcomes. There does not yet appear to be strong consensus within the literature as to which specific approaches are the most appropriate for given scenarios, or an agreed set of selection criteria. However, it is apparent that the decision on which method to adopt in a given scenario is likely to be guided, to a large extent, by the information and resource available to the practitioner. A key element is the availability of damaged state data—a topic of general concern in SHM. Methods that are based on classification outcomes will, in general, require the availability of either experimentally-obtained representations of damaged state data or, more realistically, numerical model predictions of damaged state system response. Therefore, the quality and robustness of sensor placement outcomes will therefore be a function of the information one is able or prepared to gather and include within the analysis. These concepts are explored further in Section 5.

4. Optimisation Methods

The sensor placement optimisation task (whether for SHM purposes or otherwise) is typically approached as a combinatorial optimisation problem i.e., the aim is to select some subset of size m from a discrete, finite set of size n that optimises a chosen cost function. The approaches that are available for performing combinatorial optimisation of the cost functions introduced in Section 3 are largely drawn from two categories: sequential and metaheuristic methods. Where computationally feasible, it is deemed to be good practice to compare optimisation outcomes to those of an exhaustive search.

4.1. Sequential Sensor Placement Methods

Sequential sensor placement (SSP) methods offer a set of simple to implement and computationally efficient heuristic algorithms that are appropriate to the sensor placement task. There are two principal approaches: backward sequential sensor placement (BSSP) and forward sequential sensor placement (FSSP). Both of the methods proceed along similar lines, with the BSSP algorithm starting with a large sensor set and sequentially eliminating those sensor locations that contribute least to the objective function (as in the original EI study by Kammer [6]), and FSSP beginning with a small set and adding sensor locations that offer the greatest benefit (as explored for EI in ref. [11]). These approaches are discussed in detail in refs. [4,27], with the performance of the sub-optimal sensor distributions returned by the FSSP and BSSP method shown to provide a good approximation to the performance of the optimal sensor distribution on a numerical test case.

Stephan [28] proposes a method that acts as a hybrid of the FSSP and BSSP methods. The approach is demonstrated for modal test planning on the basis of a finite element model of an aircraft. Two steps are carried out within each iteration of the algorithm. First, a new sensor location is proposed, based upon its contribution to the Fisher information matrix. Next, the levels of redundancy between the newly proposed sensor and each sensor already within the proposed distribution is evaluated, with any that fall below a pre-set threshold value eliminated. The method is shown to perform well in comparison to the standard EI method, with the clustering of sensor locations reduced.

4.2. Metaheuristic Methods

In recent years, metaheuristic (and particularly evolutionary) optimisation methods have become manifest. Metaheuristic methods are stochastic in nature and they seek to find solutions that approach the global optima. There are a wide variety of such methods, with many inspired by processes occurring in nature, such as evolution or animal behaviour. These methods are given extensive coverage in ref. [3]. The general approach adopted by metaheuristic optimisation techniques is as follows:

1. generate an initial population (the First Generation) containing randomly-generated individuals. In an SPO context, each ‘individual’ is typically taken to comprise a sensor location subset;
2. evaluate the fitness of each individual within the population against a pre-defined cost function;
3. produce a new population by (a) selecting the best-fit individuals for reproduction and (b) breeding new individuals through crossover and mutation operations; and,
4. repeat steps (2) and (3) above until some stopping criteria is reached.

4.2.1. Genetic Algorithms (GA)

GAs offer a powerful class of metaheuristic method for the solution of search and optimisation problems. They are directly inspired by the concept of natural selection with candidate solutions being encoded as gene vector, usually in the form of a binary string. GAs have been extensively used in SPO applications. A typical formulation is for the gene vector to be of length n , with each bit within the vector indicating whether or not a sensor is placed at that location.

Yao et al. [10] presented one of the earliest applications of a GA for sensor placement. Developing from the standard EI formulation proposed in ref. [6], which makes use of a backward sequential approach, the determinant of the FIM is adopted as the fitness function for the GA. The approach is applied to SPO for modal identification with two case studies considered: a space structure and a photo-voltaic (PV) array. It was found that the GA marginally outperformed the EI algorithm for an $m = 10$ and $m = 20$ measurement set case for the space structure. However, the GA was found to be significantly more costly, incurring 30 times the computational cost for the 20-sensor case. Similar results were observed for the PV array. It was noted that the GA used may converge to a local minimum, highlighting a general issue with search methods of this type. A typical approach to overcome this is to repeat the optimisation multiple times with randomisation of the starting population.

Stabb & Bleslloch [29] extended the application of GAs for SPO in a modal analysis context to two further fitness functions. The first of these is simple the MAC. The second is the cross-orthogonality matrix between the modeshapes of interest, partitioned according to the candidate sensor set and the associated Guyan-reduced mass matrix. The MAC approach was found to be quicker thanks to the lower computational cost of each evaluation, but was also found to be the less accurate of the two, as the FE modes are only truly orthogonal with respect to the mass matrix. For each measure, an error function is computed by taking the absolute value of the difference with an “ideal” matrix: the identity in the case of the cross-orthogonality matrix, and the MAC from the full modeshapes in the case of the MAC. The GA is compared with the EI and MKE methods, as well as the iterative Guyan-reduction method that is introduced in Penny et al. [12]. For a case study that is based upon an FE model of a, it was found that the best solution of all the alternatives tested was achieved using the method incorporating the Guyan-reduced mass matrix and GA optimisation, with the GA seeded using the results of a traditional EI approach. This method produces accurate modal frequencies for four out of six modes, whereas the traditional EI and KE methods performed poorly.

In one of the earliest studies on sensor placement specific to an SHM problem, Worden et al. [30] use a neural network (NN) for the location and classification of faults and a GA to determine an optimal (or near optimal) sensor configuration. The NN is trained using modeshape curvatures provided by an FE model of a cantilever plate. Faults in the plate were simulated by “removing” groups of elements in the model. The probabilities of misclassification for the different damage conditions are obtained from

the neural network, and the inverse of this measure is employed as a fitness function for the GA. A set of 20 locations were selected as a candidate sensor location set, and the GA optimisation compared to the results of two heuristic approaches: a deletion strategy and an insert strategy. The results of this preliminary study showed the GA to outperform both heuristic methods.

In Worden and Staszewski [31], a NN/GA approach was used to place sensors for the location and quantification of impacts on a composite plate. In this experimental survey, the strain data are recorded for impacts applied to a wingbox structure. An artificial neural network is trained to locate the point of impact, and a second NN employed to estimate the impact force. A GA was used to select an optimal set of three measurement locations from the candidate set of 17 used in the test, while using the estimation of the impact force provided by the neural network as the parameter to be maximised. As this parameter is dependent upon the starting conditions used to train the neural network, the training was run six times to allow for comparison between GA results and those of an exhaustive search. For each run, the GA found the same optimal set as the exhaustive search.

Coverley and Staszewski [32] presented an alternative damage location method combining classical triangulation procedures with experimental wave velocity analysis and GA optimisation. The triangulation approach assumes that three different angles for wave propagation directions from an unknown impact location are available. The distance between the sensor and the unknown impact position is calculated using the arrival times and the velocities of the propagating elastic waves. While the method was shown to perform well, a limitation of the triangulation approach is the need to estimate the velocities of the propagating elastic waves—something that is challenging in anisotropic materials. De Stefano et al. [33] adopted a trilateration approach that attempts to overcome these issues. Trilateration does not require the measurement of angles of approach, instead being based purely on the measurement of distances. The resulting system of equations is non-linear in the solution parameters, thus requiring a non-linear least-squares algorithm in order to calculate a solution. The major advantage of the trilateration method is that a single NN might be trained prior to optimisation taking place, thus offering a major reduction in computational cost in comparison to schemes where a NN must be trained at each iteration within the optimisation algorithm.

One of the key issues discussed in ref. [26] relates to appropriate methods for encoding sensor locations. The simplest way to encode an m -sensor distribution is in an n -bit binary vector, where n is the size of the candidate location set and each bit switches on if the corresponding sensor is present. However, this is suboptimal as the distribution for the sensor sets will follow a binomial distribution, and will thus be very sharply peaked at $n/2$ sensors. If n is large, randomly-generated distributions will almost certainly have close to $n/2$ sensors, regardless of the value of m . This means that crossover or mutation will usually be destructive unless protected operations are defined. The simplest alternative is to use an integer-valued GA where there are m integers, each indexing a sensor in the candidate set. Protected operators may still be convenient, or a penalty function can be used to eliminate distributions with repeats.

4.2.2. Simulated Annealing (SA)

In ref. [26], simulated annealing was one of three sensor placement methods compared using an FE model of a cantilever plate. Fault diagnosis was performed while using a neural network trained on modeshape and curvature data generated using the numerical model, resulting in the structure being classified as either faulted or not faulted. The objective function that was chosen for the SA algorithm was the probability of misclassification. The SA method was shown to perform impressively, showing a slight improvement on a comparable GA approach when 10 sensors were selected from a candidate set of 20. The SA algorithm yielded a four-sensor distribution with a 99.5% probability of correct classification, and a three-sensor distribution that was only slightly different to that found by exhaustive search.

4.2.3. Ant Colony and Bee Swarm Metaphors

The same neural network problem as in ref. [31] is investigated using a different optimisation technique in Overton and Worden [34]. Here, an ant colony metaphor was used to place sensors based upon the fitness functions that are generated by the neural network, and the results compared with those from the GA and exhaustive search. For the selection of a three-sensor distribution from a set of 17, the ant colony algorithm is shown to be faster than the GA while producing a solution that is amenable to that from the exhaustive search. The algorithm also appeared to outperform the GA when a six-sensor distribution was sought, although entirely fair comparison between the approaches was not possible. As the size of the sensor network increases, it appears that the selected distributions become less intuitive.

In common with other NN-based approaches, it is noted that the objective function adopted in ref. [34] is a random variable conditioned on the initial weights given to the network. Scott and Worden [35] describe the development of a bee swarm algorithm, again via application to the NN-based objective function proposed in ref. [31]. The bee swarm algorithm was shown to perform well for an initial, constrained set of neural network parameters. However, the key development was to extend the algorithm in order to include optimisation of the neural network parameters alongside the sensor distributions. In this set up the bee swarm algorithm was shown to significantly improve accuracy, outperforming both the ant colony algorithm and the genetic algorithm in cases where performance could be directly compared. It was noted that the nature of the NN-based objective function was well-suited to the bee swarm algorithm, which only requires there to be some similarities—however vague—between “good” solutions in order to substantially outperform a random search.

4.2.4. Mixed Variable Programming (MVP)

In Beal et al. [36], optimal sensor placement is formulated as a mixed variable programming (MVP) problem. As observed in ref. [26], a typical formulation for the SPO process is to treat each sensor location as a binary variable (the sensor is either placed at the given location, or not), with the resulting optimisation being a binary mixed integer non-linear programming problem. A difficulty that arises from this mixed integer formulation is that it scales poorly, becoming computationally demanding as the candidate sensor set becomes large. The proposed alternative is a mixed variable programming scheme, where each sensor is assigned a categorical variable from a predefined list indicating where it is to be located. The MVP framework is sufficiently general that it may be applied a broad set of objective functions. The objective function used for demonstration in ref. [36] is the relative error between the stiffness changes indicated by the full and reduced measurement sets. The approach is demonstrated using a numerical model of a 20-DOF system. Given that damage had occurred at specified masses, optimal configurations containing three-, four-, and five-sensors were found for three damage cases. A convergence analysis is presented that indicates strong results for optimisation problems with one continuous variable, indicating that the approach is sufficiently general for a variety of SPO problems.

4.3. Summary

Given the extensive literature in this area, it has become increasingly challenging for authors in this area to demonstrate true novelty of approach. This issue is exacerbated in cases where there is no principled comparison either to other methods or to best possible outcomes. The best possible outcome might be interpreted either as the outcome achievable if all data are available (i.e., the full sensor set may be used); or, the best possible outcome achievable for a given sensor subset size, evaluated via exhaustive search. In either interpretation, the best possible performance represents an upper bound on system performance.

Also noted is that the connection between the problem to be solved (minimisation of a particular cost function relating to sensor placement) and the optimisation method chosen is given little explicit

consideration in many of the papers that appear in the SPO literature. The “no free lunch” theorems popularised by Wolpert and Macready [37] establish that for any given optimisation algorithm, over-performance for one class of problems will be offset by under-performance for another class. Thus, for uniformly distributed problem classes (i.e., the likelihood of being faced with a “simple” optimisation task is the same as being presented with a “hard” task), the performance of any given optimisation algorithm will be equal to that of any other when averaged across all classes. It is only by making use of *a priori* understanding of the particular problem at hand that overall improved performance can be achieved.

5. Emerging Trends and Future Directions

This section seeks to identify research directions that have the potential to make a substantial contribution to the problem of system design for SHM. It is apparent from the quantity of papers available in the literature, of which only a subset are reviewed in Sections 3 and 4, where sensor placement optimisation is, in some senses, a mature field. Optimisation methods applicable to the combinatorial subset selection task have, in particular, been extensively explored. While there may still be some scope for tailoring optimisation methods to particular approaches in the pursuit of computational efficiency, there appears general agreement that for any given appropriately-posed problem, an acceptable solution will be provided by the methods that are covered in Section 4. There is arguably more scope for novelty in regard to cost functions, particularly with regard to tailoring cost functions to specific SHM systems. However, the avenue for real progress appears to be in directly addressing the trade offs that are required to make a decision on whether or not to implement an SHM system. There is agreement here with the perspectives expressed in ... and more recently of ref. [38] that the focus of SHM research should—at this stage in its development—be on moving towards industrial deployment of monitoring systems.

Key among the trends identified and discussed are applications of methods that are drawn from decision theory (Section 5.1) and the quantification of overall SHM system value (Section 5.2). A comparatively recent trend is to seek to make the sensor network robust to confounding factors, including: (1) benign effects driven by Environmental and Operational Variables (EOVs); (2) uncertainties and errors in any contributing mathematical or numerical models; and, (3) sensor and sensor network failure. The emerging literature in each area is briefly reviewed in Sections 5.3 and 5.4. Finally, methods that consider robustness to sensor or sensor network failure are considered (Section 5.5), as are the opportunities for gathering full-field, experimental candidate set data via laser vibrometry (Section 5.6). Several of the reviewed papers do not explicitly cover sensor placement optimisation for SHM, but introduce ideas or methods that are relevant to further development in this field.

5.1. From Classifiers to Decisions

A key emerging trend within the research literature is to move from the optimisation of classification outcomes to the optimisation of decision outcomes. Flynn and Todd [39] present perhaps the most influential recent study in this field, setting out a Bayesian risk minimisation framework for optimal sensor placement. This framework enables the costs associated with different actions to be included in the decision-making process, enabling a design that minimises risk under uncertainty to be found. Here, design, is taken to incorporate sensor placement, feature selection, and the setting of localised detection thresholds. This essentially casts damage identification as a binary (or more generally M -ary) hypothesis testing problem, weighted by the cost of making a correct/incorrect decision. The framework proposed is general and, in principle, could be applied to multisite damage cases with costs that are associated with all outcomes. However, through a series of assumptions the approach is simplified to the consideration of binary classifiers that are associated with local health states, with the structure discretised into K spatial regions, indexed by $k = 1 \dots K$. Figure 3 shows a schematic illustration of a discretised structure. For each region, a local detection

threshold $T[k] > \gamma[k]P(h_{k0})/P(h_{k1})$ is established, where $\gamma[k]$ represents the classification cost ratio, h_{kj} represents the true local damage state (with $j = 0$ and $j = 1$ denotes undamaged and damaged states, respectively), and $P(h_{kj})$ represents prior probabilities on the local damage states.

Figure 3. Schematic example of a discretised structure.

Given these assumptions, the global probability of detection P_D is given by

$$P_D = \sum_{k=1}^K \frac{P(d_{k1}|h_{k1})P(h_{k1})}{\sum_{k=1}^K P(h_{k1})} \quad (17)$$

where d_{kj} represents the event of deciding that h_{kj} is the local damage state in region k and the local detection rate is given by,

$$P(d_{k1}|h_{k0}) = P_r(T[k] > \gamma[k]|h_{k1}) \quad (18)$$

Similarly, the global probability of false alarm rate P_{FA} is given by,

$$P_{FA} = \sum_{k=1}^K \frac{P(d_{k1}|h_{k0})P(h_{k0})}{\sum_{k=1}^K P(h_{k1})} \quad (19)$$

with the local false alarm rate given by,

$$P(d_{k1}|h_{k1}) = P_r(T[k] > \gamma[k]|h_{k0}) \quad (20)$$

While the particular case study in ref. [39] focuses on active sensing, the framework presented is appropriate for general application. The advantage of the assumptions made is that closed-form analytical expressions may be derived for P_D and P_{FA} as a function of the costs, prior probabilities of damage and local threshold damage threshold values. The optimisation task is then to maximise P_D for a given allowable false alarm rate P_{FA} ; or alternatively, minimise P_{FA} for a given target P_D . The key outstanding question that remains is how best to establish the feature distributions that are associated with each health state of interest.

5.2. Quantifying the Value of SHM Systems

Going further, Thöns [40] presents a study considering the value of an SHM system as the quantity of primary concern. A framework for quantification of the value of SHM in the context of the structural risk and integrity management is presented. This approach seeks to directly address the core decision to be made by an asset manager: can the installation of the proposed system be justified? The approach adopted builds, as in ref. [39], on the notion of expected utilities associated with available decisions. The framework is demonstrated for a structural system undergoing a process of fatigue damage, with decision processes modelled in detail along with associated uncertainties. It is demonstrated that the value of information provided by an SHM system varies significantly between contexts, with the relative value being particularly strongly influenced by factors including: modelling of brittle and ductile behaviour of components, adopted fatigue damage thresholds and whether the system is able to identify applied loads. While the SHM system design is not explicitly considered in ref. [40], the

study presents the possibility of defining new cost functions that are based on the value of information provided by the system, directly informing decision making.

5.3. Robustness to Benign Effects

The systems of interest are typically exposed to significant operational and environmental effects, and it is well established that the features that are typically adopted for SHM purposes also demonstrate significant sensitivity to these benign effects. Among the most highly-cited illustrations of the impact of temperature effects is the Z-24 Bridge dataset gathered by Peeters and De Roeck [41]. In addition to environmental effects, load variation is a particular concern for structures that exhibit excitation-dependent nonlinear behaviour. Accounting for benign effects is an area of active research in its own right with concepts, such as latent variable models [42] and cointegration [43] having been demonstrated to work to great effect. Nonetheless, residual effects remain and should ideally be considered during SHM system design.

Li et al. [44] propose a load-dependent sensor placement method that is based on the EI algorithm. The adaptation made is to develop an objective function that minimises the Euclidean distance between a modal estimator based on a selected sensor subset and an ideal estimator that uses all available measurement locations in a manner similar to that proposed in ref. [26]. The considered load variation refers to the location of excitation, with the magnitude of loading not varied. A case study on a six-storey truss structure with a distribution comprising six sensors being selected from 12 candidate locations is used to demonstrate that the location at which the load is applied has a marked influence on the selected sensor layout. While the robustness of the sensor network to variations in load location was not explicitly considered, this would appear to be a straightforward progression from the work presented.

5.4. Robustness to Modelling and Prediction Errors

The aforementioned study by Kammer [8] presented an early effort to assess the effect of model error on sensor placement outcomes. In recent years, this has been developed much further with the aim of accounting for errors in the underlying model used for SHM system design. Papadimitriou and Lombaert [45] consider the effect of spatially-correlated prediction errors on sensor placement outcomes. Spatially-correlated errors are considered to be important, as they influence the degree of redundancy observed in the information content of adjacent sensors. A sequential sensor placement algorithm is adopted for optimisation, and it is demonstrated that algorithms of this type remained both accurate and computationally efficient in the face of spatially-correlated prediction errors.

Castro-Triguero et al. [46] examined the influence of parametric modelling uncertainties on optimal sensor placement for the purposes of modal analysis. The effect of uncertainty is demonstrated for four SPO approaches applied to a truss bridge structure on the basis of FE model predictions. A combination of material (Young's modulus, density) and geometric (cross-sectional dimension) model parameters were varied, with uncertainties being propagated through to model predictions via Monte Carlo sampling. In addition, noise on measured outputs was simulated. The considered numerical case study illustrated that the effect of parametric uncertainties on the returned sensor configuration was significant, although some "key" sensor locations frequently recurred within the selected sensor sets.

5.5. Robustness to Sensor and Sensor Network Failure

The work presented in ref. [31] is extended in ref. [47] to cover fail-safe sensor placements. The focus of the study is the selection of sensor sets whose subsets also offer a low rate of detection error (equivalently a high probability of correct classification P_{corr}). A fail-safe fitness function is introduced that operates as follows: given a proposed mother distribution comprising N sensors, all $N - 1$ child distributions are generated and their classification performance evaluated. The value adopted for the performance of the mother set is the worst P_{corr} performance of its child distributions.

The approach is illustrated for the case of impact magnitude estimation on a composite plate using time-varying strain data gathered from a network of 17 piezoceramic sensors. A sensor subset size of three was selected for illustration, making it feasible to compare the results of a GA to the optimal outcome provided by an exhaustive search. The key observation made is that, for the particular case studied, there was a trade off between achieving an optimal mother distribution, and one that is robust to the failure of one sensor. A further observation made is that the optimal sensor distribution returned by the GA displayed marked sensitivity to statistical variations in the probability of detection error, highlighting the value in pursuing robust optimisation methods. While demonstrated for strain data and the case that only a single sensor fails, the proposed fail-safe approach was deemed sufficient generally for extension to other SHM problems to be straightforward.

Kullaa [48] presented a study on methods for distinguishing between structural damage, EOVs and sensor faults. The approach is based on the fact that sensor faults are local, while structural damage and EOVs are global. EOVs are dealt with by including these effects within the training set, in a manner consistent with the supervised learning approaches discussed in section 5.3. The study is structured such that each sensor yields a single test statistic. A Gaussian Process (GP) regression model is trained across all sensor locations such that the output of individual sensors may be estimated using the output from all others. A hypothesis test (in this case the generalised likelihood ratio test (GLRT)) is then applied to identify whether the measured output of a given sensor location deviates significantly from that predicted by the GP, enabling sensor faults to be identified. While sensor optimisation is not explicitly considered, the work presents a basis for building networks that are robust to sensor failure.

More recently, ref. [49] considers fault-tolerant wireless sensor network (WSN) design for SHM. The adoption of a wireless network architecture is an attractive development for SHM systems, offering simplified deployment that obviates the need for extensive cabling. However, wireless architectures introduce an additional set of failure modes (for example, communication errors and unstable connectivity) in addition to failure of the sensor itself. They also introduce an additional energy constraint that is not considered for wired systems. The objective of the work presented in ref. [49] is to maintain a given level of fault tolerance through robustly-optimal system design. This is achieved by implementing an algorithm that identifies optimal repairing points within sensor clusters and places backup sensors at those locations. In a further adaptation, the adopted SHM algorithm makes use of decentralised computing, with the objective that the likelihood of the network remaining connected in the event of system fault is maximised.

Finally, Yi et al. [50] present a survey of recent literature in the field of sensor validation methodologies for structural health monitoring.

5.6. Scanning Laser Vibrometry

Two primary sources of candidate sensor location information have been proposed in the literature: the predictions of a (typically FE) model; and experimentally-measured data from an extensive sensor array. The advent and increasing availability of Scanning Laser Vibrometry (SLV) presents the possibility of achieving full field vibration data acquisition experimentally, removing any issues that may arise through test-model bias. The fact this may be done in a non-contacting fashion, removing structural effects such as mass loading during acquisition of candidate set data, is a further boon. Staszewski et al. [51] demonstrated the potential of scanning laser vibrometry for Lamb wave sensing. More recently Marks et al. [52] demonstrated optimisation of sensor locations for Lamb wave based damage detection using a GA allied to SLV data. However, there remain comparatively few studies that explore the use of full-field measurement techniques as a source of candidate set data.

6. Concluding Comments

This paper presents a survey and comparison of SPO approaches applicable to SHM task. The need for sensor placement optimisation techniques specific to SHM has long been recognised and promising approaches have emerged. A great number of techniques developed for other structural dynamics

applications (modal analysis, system identification, and model calibration) are also of use to the SHM practitioner. Future methods are likely to be specific to the SHM methodology employed and should be developed alongside them. The central focus of this review was to highlight emerging trends specific to SHM system development, and these are briefly summarised below.

While there has been a great deal of focus on optimisation algorithms, it is the authors' opinion that as the optimisation cost is incurred at the design stage rather than online, the costs that are associated with this step are arguably of minor importance in comparison to the overall cost of developing, implementing, and maintaining an SHM system. In terms of the development of new optimisation algorithms, two recommendations present themselves. The first is that prior to presenting the outcomes of optimising an SHM system design, it is generally desirable that the best possible performance of the system (as defined in Section 4.3) be evaluated to demonstrate whether an acceptable level of performance is achievable. Secondly, it is desirable that the performance of an "optimal" sensor distribution returned via a newly-proposed algorithm should be compared with the results of an exhaustive search wherever feasible. This is admittedly limited to cases with low numbers of sensors in the candidate and measurement sets.

There remains scope for being more ambitious in the damage-sensitive features adopted. This ambition has two strands. The first is that there remains a focus in large parts of the literature on modal features alone, although a smaller number of the reviewed studies considered other options based on strain field measurements and time of flight of wave packets for impact detection, for example. Methods appropriate to general feature sets including multiple sensor modalities—the domain of data fusion—remain an emerging topic. The second area for greater ambition is in incorporating feature selection within the sensor placement task i.e the selection of maximally sensitive feature sets within the data gathered from a given sensor distribution). A common approach in the literature is to set the features to be returned by a sensor location at an early stage (e.g., a vector of modeshape elements). While removing such constraints may inevitably, and perhaps dramatically, increase the computational expense associated with the optimisation task, it may extend the upper bound on "best possible" performance.

As highlighted in Section 5.4, there have been some efforts towards developing SPO methods that are robust to model uncertainty and experimental variability, but there remains scope to go much further with this in order to increase confidence in the SHM system design process. There also remains a question mark over how to overcome the difficulty of sourcing damaged state data with which to train SHM classifiers. The sources of data used for SHM system optimisation fall broadly into two categories: the predictions of a numerical model of a given structure, or experimental data gathered from a structure in situ. The former allows for consideration of the effects of damage, but legitimate questions may be posed around the required accuracy of damaged-state model predictions. The latter avoids the issue of test-model bias but will typically be limited by the lack of system-level damaged-state data (although advances are being made in the use of Population-Based methods that attempt to address this [53]). The overwhelming majority of studies reviewed focus purely on development and evaluation using numerical data; or alternatively demonstrated these steps using purely experimental data. It is noted that there remains a distinct shortage of work concentrating on the practical, experimental evaluation of systems developed using *a priori* physics-based modelling on realistic structures. Finally, there appears to be broad scope for the further development of methods that focus on optimisation of SHM decision outcomes, as opposed to classification outcomes alone.

Author Contributions: R.B.: Conceptualization, writing—original draft preparation; K.W.: Conceptualization, writing—review and editing. All authors have read and agreed to the published version of the manuscript.

Funding: This work was supported by the Engineering and Physical Sciences Research Council [Grant numbers EP/J013714/1 and EP/N010884/1].

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

The following abbreviations are used in this manuscript:

SHM	Structural health monitoring
SPO	Sensor placement optimisation
DOFs	Degrees of freedom
FEA	Finite element analysis
EO	Evolutionary optimisation
FIM	Fisher information matrix
EOV	Experimental and operational variation
MAC	Modal assurance criterion
NMSE	Normalised mean squared error
GA	Genetic algorithm
EI	Effective independence
SSP	Sequential sensor placement
ROC	Receiver operating characteristic

References

1. Barthorpe, R.J.; Worden, K. Sensor Placement Optimization. In *Encyclopedia of Structural Health Monitoring*; Boller, C., Chang, F., Fujino, Y., Eds.; John Wiley & Sons: Chichester, UK, 2009. [\[CrossRef\]](#)
2. Yi, T.H.; Li, H.N. Methodology Developments in Sensor Placement for Health Monitoring of Civil Infrastructures. *Int. J. Distrib. Sens. Netw.* **2012**, *8*. [\[CrossRef\]](#)
3. Ostachowicz, W.; Soman, R.; Malinowski, P. Optimization of sensor placement for structural health monitoring: A review. *Struct. Health Monit.* **2019**, *18*, 963–988. [\[CrossRef\]](#)
4. Papadimitriou, C. Optimal sensor placement methodology for parametric identification of structural systems. *J. Sound Vib.* **2004**, *278*, 923–947. [\[CrossRef\]](#)
5. Udwadia, F. Methodology for Optimum Sensor Locations for Parameter Identification in Dynamic Systems. *ASCE J. Eng. Mech.* **1994**, *120*. [\[CrossRef\]](#)
6. Kammer, D. Sensor placement for on-orbit modal identification and correlation of large space structures. *J. Guid. Control. Dyn.* **1991**, *14*, 251–259. [\[CrossRef\]](#)
7. Shah, P.; Udwadia, F. A Methodology for Optimal Sensor Locations for Identification of Dynamic Systems. *J. Appl. Mech.* **1978**, *45*, 188–196. [\[CrossRef\]](#)
8. Kammer, D. Effect of model error on sensor placement for on-orbit modal identification of large space structures. *J. Guid. Control. Dyn.* **1992**, *15*, 334–341. [\[CrossRef\]](#)
9. Kammer, D. Effects of noise on sensor placement for on-orbit modal identification of large space structures. *J. Dyn. Syst. Meas. Control. Trans. ASME* **1992**, *114*, 436–443. [\[CrossRef\]](#)
10. Yao, L.; Sethares, W.; Kammer, D. Sensor placement for on-orbit modal identification via a genetic algorithm. *AIAA J.* **1993**, *31*, 1922–1928. [\[CrossRef\]](#)
11. Kammer, D. Sensor set expansion for modal vibration testing. *Mech. Syst. Signal Process.* **2005**, *19*, 700–713. [\[CrossRef\]](#)
12. Penny, J.; Friswell, M.; Garvey, S. Automatic choice of measurement locations for dynamic testing. *AIAA J.* **1994**, *32*, 407–414. [\[CrossRef\]](#)
13. Salama, M.; Rose, T.; Garba, J. Optimal placement of excitations and sensors for verification of large dynamical systems. In Proceedings of the 28th Structures, Structural Dynamics and Materials Conference, Monterey, CA, USA, 6–8 April 1987. [\[CrossRef\]](#)
14. Li, D.S.; Li, H.N.; Fritzen, C.P. The connection between effective independence and modal kinetic energy methods for sensor placement. *J. Sound Vib.* **2007**, *305*, 945–955. [\[CrossRef\]](#)
15. Imamovic, N.; Ewins, D. Optimization of excitation DOF selection for modal tests. In Proceedings of the 15th International Modal Analysis Conference, Orlando, FL, USA, 3–6 February 1997.
16. Imamovic, N. Validation of Large Structural Dynamics Models Using Modal Test Data. Ph.D. Thesis, Imperial College London, London, UK, 1998.

17. Larson, C.B.; Zimmerman, D.C.; Marek, E.L. A comparison of modal test planning techniques: Excitation and sensor placement using the NASA 8-bay truss. In Proceedings of the 12th International Modal Analysis Conference, Honolulu, HI, USA, 31 January–3 February 1994; pp. 205–211.
18. Shi, Z.; Law, S.; Zhang, L. Optimum sensor placement for structural damage detection. *J. Eng. Mech.* **2000**, *126*, 1173–1179. [[CrossRef](#)]
19. Guo, H.; Zhang, L.; Zhang, L.; Zhou, J. Optimal placement of sensors for structural health monitoring using improved genetic algorithms. *Smart Mater. Struct.* **2004**, *13*, 528–534. [[CrossRef](#)]
20. Hemez, F.; Farhat, C. An energy based optimum sensor placement criterion and its application to structural damage detection. In Proceedings of the SPIE International Society for Optical Engineering, Bergen, Norway, 13–15 June 1994; pp. 1568–1568.
21. Said, W.M.; Staszewski, W.J. Optimal sensor location for damage detection using mutual information. In Proceedings of the 11th International Conference on Adaptive Structures and Technologies (ICAST), Nagoya, Japan, 23–26 October 2000; pp. 428–435.
22. Trendafilova, I.; Heylen, W.; Van Brussel, H. Measurement point selection in damage detection using the mutual information concept. *Smart Mater. Struct.* **2001**, *10*, 528–533. [[CrossRef](#)]
23. Papadimitriou, C.; Beck, J.; Au, S.K. Entropy-based optimal sensor location for structural model updating. *Jvc/J. Vib. Control.* **2000**, *6*, 781–800. [[CrossRef](#)]
24. Tharwat, A. Classification assessment methods. *Appl. Comput. Inform.* **2018**. [[CrossRef](#)]
25. Fawcett, T. An introduction to ROC analysis. *Pattern Recognit. Lett.* **2006**, *27*, 861–874. [[CrossRef](#)]
26. Worden, K.; Burrows, A. Optimal sensor placement for fault detection. *Eng. Struct.* **2001**, *23*, 885–901. [[CrossRef](#)]
27. Papadimitriou, C. Pareto optimal sensor locations for structural identification. *Comput. Methods Appl. Mech. Eng.* **2005**, *194*, 1655–1673. [[CrossRef](#)]
28. Stephan, C. Sensor placement for modal identification. *Mech. Syst. Signal Process.* **2012**, *27*, 461–470. [[CrossRef](#)]
29. Stabb, M.; Bleslo, P. A genetic algorithm for optimally selecting accelerometer locations. In Proceedings of the 13th International Modal Analysis Conference, Nashville, TN, USA, 13–16 February 1995; Volume 2460, p. 1530.
30. Worden, K.; Burrows, A.; Tomlinson, G. A combined neural and genetic approach to sensor placement. In Proceedings of the 15th International Modal Analysis Conference, Nashville, TN, USA, 12–15 February 1995; pp. 1727–1736.
31. Worden, K.; Staszewski, W. Impact location and quantification on a composite panel using neural networks and a genetic algorithm. *Strain* **2000**, *36*, 61–70. [[CrossRef](#)]
32. Coverley, P.; Staszewski, W. Impact damage location in composite structures using optimized sensor triangulation procedure. *Smart Mater. Struct.* **2003**, *12*, 795. [[CrossRef](#)]
33. De Stefano, M.; Gherlone, M.; Mattone, M.; Di Sciuva, M.; Worden, K. Optimum sensor placement for impact location using trilateration. *Strain* **2015**, *51*, 89–100. [[CrossRef](#)]
34. Overton, G.; Worden, K. Sensor optimisation using an ant colony metaphor. *Strain* **2004**, *40*, 59–65. [[CrossRef](#)]
35. Scott, M.; Worden, K. A Bee Swarm Algorithm for Optimising Sensor Distributions for Impact Detection on a Composite Panel. *Strain* **2015**, *51*, 147–155. [[CrossRef](#)]
36. Beal, J.M.; Shukla, A.; Brezhneva, O.A.; Abramson, M.A. Optimal sensor placement for enhancing sensitivity to change in stiffness for structural health monitoring. *Optim. Eng.* **2007**, *9*, 119–142. [[CrossRef](#)]
37. Wolpert, D.H.; Macready, W.G. No free lunch theorems for optimization. *IEEE Trans. Evol. Comput.* **1997**, *1*, 67–82. [[CrossRef](#)]
38. Cawley, P. Structural health monitoring: Closing the gap between research and industrial deployment. *Struct. Health Monit.* **2018**, *17*, 1225–1244. [[CrossRef](#)]
39. Flynn, E.; Todd, M. A Bayesian approach to optimal sensor placement for structural health monitoring with application to active sensing. *Mech. Syst. Signal Process.* **2010**, *24*, 891–903. [[CrossRef](#)]
40. Thöns, S. On the Value of Monitoring Information for the Structural Integrity and Risk Management. *Comput. Civ. Infrastruct. Eng.* **2018**, *33*, 79–94. [[CrossRef](#)]
41. Peeters, B.; De Roeck, G. One-year monitoring of the Z24-Bridge: Environmental effects versus damage events. *Earthq. Eng. Struct. Dyn.* **2001**, *30*, 149–171. [[CrossRef](#)]

42. Kullaa, J. Eliminating Environmental or Operational Influences in Structural Health Monitoring using the Missing Data Analysis. *J. Intell. Mater. Syst. Struct.* **2009**, *20*, 1381–1390. [[CrossRef](#)]
43. Cross, E.J.; Worden, K.; Chen, Q. Cointegration: A novel approach for the removal of environmental trends in structural health monitoring data. *Proc. R. Soc. A* **2011**, *467*, 2712–2732. [[CrossRef](#)]
44. Li, D.S.; Li, H.N.; Fritzen, C.P. Load dependent sensor placement method: Theory and experimental validation. *Mech. Syst. Signal Process.* **2012**, *31*, 217–227. [[CrossRef](#)]
45. Papadimitriou, C.; Lombaert, G. The effect of prediction error correlation on optimal sensor placement in structural dynamics. *Mech. Syst. Signal Process.* **2012**, *28*, 105–127. [[CrossRef](#)]
46. Castro-Triguero, R.; Murugan, S.; Gallego, R.; Friswell, M.I. Robustness of optimal sensor placement under parametric uncertainty. *Mech. Syst. Signal Process.* **2013**, *41*, 268–287. [[CrossRef](#)]
47. Side, S.; Staszewski, W.; Wardle, R.; Worden, K. Fail-safe sensor distributions for damage detection. In Proceedings of the DAMAS'97 International Workshop on Damage Assessment using Advanced Signal Processing Procedures, Sheffield, UK, 30 June–2 July 1997; pp. 135–146.
48. Kullaa, J. Distinguishing between sensor fault, structural damage, and environmental or operational effects in structural health monitoring. *Mech. Syst. Signal Process.* **2011**, *25*, 2976–2989. [[CrossRef](#)]
49. Bhuiyan, M.; Wang, G.; Cao, J.; Wu, J. Deploying Wireless Sensor Networks with Fault-Tolerance for Structural Health Monitoring. *IEEE Trans. Comput.* **2015**, *64*, 382–395. [[CrossRef](#)]
50. Yi, T.H.; Huang, H.B.; Li, H.N. Development of sensor validation methodologies for structural health monitoring: A comprehensive review. *Measurement* **2017**, *109*, 200–214. [[CrossRef](#)]
51. Staszewski, W.; Lee, B.; Mallet, L.; Scarpa, F. Structural health monitoring using scanning laser vibrometry: I. Lamb wave sensing. *Smart Mater. Struct.* **2004**, *13*, 251. [[CrossRef](#)]
52. Marks, R.; Clarke, A.; Featherston, C.A.; Pullin, R. Optimization of acousto-ultrasonic sensor networks using genetic algorithms based on experimental and numerical data sets. *Int. J. Distrib. Sens. Netw.* **2017**, *13*. [[CrossRef](#)]
53. Gardner, P.; Liu, X.; Worden, K. On the application of domain adaptation in structural health monitoring. *Mech. Syst. Signal Process.* **2020**, *138*, 106550. [[CrossRef](#)]

© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).