



This is a repository copy of *A probabilistic risk-based decision framework for structural health monitoring*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/167396/>

Version: Published Version

Article:

Hughes, A.J. orcid.org/0000-0002-9692-9070, Barthorpe, R.J. orcid.org/0000-0002-6645-8482, Dervilis, N. orcid.org/0000-0002-5712-7323 et al. (2 more authors) (2021) A probabilistic risk-based decision framework for structural health monitoring. *Mechanical Systems and Signal Processing*, 150. 107339. ISSN 0888-3270

<https://doi.org/10.1016/j.ymssp.2020.107339>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:
<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>



A probabilistic risk-based decision framework for structural health monitoring

A.J. Hughes^{a,*}, R.J. Barthorpe^a, N. Dervilis^a, C.R. Farrar^b, K. Worden^a

^aDynamics Research Group, Department of Mechanical Engineering, University of Sheffield Sheffield S1 3JD, UK

^bEngineering Institute, MS T-001, Los Alamos National Laboratory, Los Alamos, NM 87545, USA

ARTICLE INFO

Article history:

Received 1 April 2020

Received in revised form 28 July 2020

Accepted 1 October 2020

Keywords:

Structural health monitoring
Probabilistic risk assessment
Probabilistic graphical models
Decision-making

ABSTRACT

Obtaining the ability to make informed decisions regarding the operation and maintenance of structures, provides a major incentive for the implementation of structural health monitoring (SHM) systems. Probabilistic risk assessment (PRA) is an established methodology that allows engineers to make risk-informed decisions regarding the design and operation of safety-critical and high-value assets in industries such as nuclear and aerospace. The current paper aims to formulate a risk-based decision framework for structural health monitoring that combines elements of PRA with the existing SHM paradigm. As an apt tool for reasoning and decision-making under uncertainty, probabilistic graphical models serve as the foundation of the framework. The framework involves modelling failure modes of structures as Bayesian network representations of fault trees and then assigning costs or utilities to the failure events. The fault trees allow for information to pass from probabilistic classifiers to influence diagram representations of decision processes whilst also providing nodes within the graphical model that may be queried to obtain marginal probability distributions over local damage states within a structure. Optimal courses of action for structures are selected by determining the strategies that maximise expected utility. The risk-based framework is demonstrated on a realistic truss-like structure and supported by experimental data. Finally, a discussion of the risk-based approach is made and further challenges pertaining to decision-making processes in the context of SHM are identified.

© 2020 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The field of engineering known as *Structural Health Monitoring* (SHM) concerns the development and implementation of data acquisition and processing systems for the purpose of damage detection in aerospace, civil or mechanical infrastructure [1]. A prime motivation for the use of SHM systems is to acquire the ability to make informed decisions regarding the operation and management of structures so as to improve safety and/or reduce costs. In the context of SHM, an agent tasked with making decisions is required to specify action policies that are robust to uncertainties that arise as a result of having imperfect information regarding the damage state of a structure. In addition, it may be desirable to manage a population of structures simultaneously; further complicating the decision problem. The task of decision-making for SHM is complex and highly involved and therefore demands a thorough and systematic approach.

* Corresponding author.

E-mail address: ajhughes2@sheffield.ac.uk (A.J. Hughes).

For many years, *Probabilistic Risk Assessment* (PRA) has been used in industries such as aerospace and nuclear for making decisions regarding the design and operation of safety-critical assets, such as nuclear power plants [2] and reusable space vehicles, like the space shuttle [3]. PRA provides a rigorous and structured methodology for identifying possible adverse events associated with the operation of a system and subsequently quantifying the respective probabilities of occurrence and the severity of the consequences.

Thus far, the majority of research in the field of SHM has been focussed on the identification, localisation and classification of damage. There have been fewer attempts to address decision-making processes and to incorporate risk into SHM problems. A static probabilistic risk-based approach to SHM was applied to a simulated truss structure in [4], although no attempt was made to forecast structural degradation. Flynn and Todd successfully applied a Bayes risk approach to the decision problem of sensor placement for an SHM system on square, gusset and T-shaped plates in [5]. The approach considered the risk of false positives and false negatives of damage identification in discrete regions of the plates. Cost-informed decision-making for miter gates was demonstrated in [6]; this involved using a Bayesian neural network trained on a finite element model to infer damage and forecast performance using a transition matrix. An approach proposed in [7], facilitates cost-efficient reliability-based maintenance. As the latter is a reliability-based approach rather than a risk-based approach, the costs of failure events and maintenance are not explicitly modelled. Hence, whilst the maintenance strategies developed may be cost-efficient for given safety parameters, they are not necessarily cost-optimal. There has been some research into the risk-based operation and maintenance of structures and components. Nielsen details a risk-based approach that is utilised for the operation and maintenance of off-shore wind turbines in [8], using probabilistic graphical modelling. Similarly, Hovgaard and Brincker provide a case study demonstrating a risk-based approach to the monitoring and maintenance of a finite element model of a wind turbine tower experiencing circumferential cracking in [9]. In [10] a continuous-state partially-observable Markov decision process (POMDP) was demonstrated on artificial data for maintenance planning on a deteriorating bridge. Dynamic Bayesian networks are also employed in [11] for the diagnosis and prognosis of the structural health of an aircraft wing; this includes the probabilistic temporal modelling and prediction of crack growth.

The current paper aims to address the lack of a generalised framework for conducting risk-based monitoring of structures at the full-system scale by augmenting the current SHM paradigm with practices employed in probabilistic risk assessment and thereby facilitating the decision-making processes that motivate the implementation of SHM systems. An overview is given of the current paradigms for conducting PRA and SHM as outlined in the literature. Also provided is background theory regarding the key technologies required for mapping PRA onto SHM; namely, probabilistic graphical models in the form of Bayesian networks and influence diagrams. A notation is established before the augmented risk-based paradigm for SHM is detailed. Finally, a discussion around the framework is made and further challenges in the SHM decision-process are identified.

2. Structural Health Monitoring

Structural health monitoring involves implementing damage identification strategies to determine the health state of a structure throughout its operational lifetime. Statistical pattern recognition (SPR) offers a natural approach to SHM as it allows the uncertainties inherent in engineering problems to be dealt with robustly. It is for this reason that the SPR approach has been the focus of much research over the past three decades. The established SPR paradigm for an SHM system is composed of four procedures [1]:

1. Operational Evaluation.
2. Data acquisition.
3. Feature selection.
4. Statistical modelling for feature discrimination.

2.1. Operational evaluation

Operational evaluation seeks to answer several questions concerning the implementation of an SHM system, specifically:

- What is the justification (safety and/or economic) for implementing an SHM system?
- How is damage defined for the system and what are the critical damage states?
- What are the environmental and operational conditions that the monitoring system is required to perform under?
- How does the operational environment limit data acquisition?

For an SHM system to be successfully developed and implemented, a substantial amount of information must be collected during the operational evaluation process, and significant effort may be necessary to obtain an adequate amount that is quantified in sufficient detail. Examples of required information include: monetary cost and reliability of the proposed SHM system; pertinent damage states of the structure and the thresholds at which they can be deemed to have occurred, and the temperature and load variations experienced by the structure during operation.

2.2. Data acquisition

The data acquisition process is informed by the operational evaluation. The process aims to finalise the types, number and locations of sensors to be used in the SHM system. The data acquisition, storage and transmittal hardware must also be selected. The process is constrained by both economic restrictions and the limitations enforced by the expected environmental conditions; for this reason, the data acquisition process is context dependent and relies heavily on the information gathered during the operational evaluation stage.

2.3. Feature selection

Once the data have been acquired, a set of features must be constructed that indicate whether or not there is damage present in the structure. This procedure often involves processing the data acquired from the structure; common practices include domain transformation, dimensionality reduction, and normalisation [1].

2.4. Statistical modelling for feature discrimination

Statistical models must be developed to exploit the discrepancy between features that indicate differing damage states. The degree of knowledge regarding the damage state obtained from an SHM system is highly dependent on the statistical model employed and can be evaluated in terms of Rytter's Hierarchy [12]:

1. Is there damage in the system?
2. Where is the damage located?
3. What type of damage is present?
4. How severe is the damage?
5. How much useful life remains?

Whilst Rytter's hierarchy, in itself, does not lead to decisions being made, as SHM systems progress up the hierarchy, the information they yield becomes increasingly useful to agents tasked with deciding upon a course of action for a structure. Within the field of decision theory, cost and utility are metrics used ubiquitously for the comparison of courses of action and their consequences. By combining the cost/utility of a given consequence with the respective likelihood, one can arrive at the notion of *risk*.

3. Probabilistic Risk Assessment

Probabilistic risk assessment (PRA) is a method that is widely used for evaluating risks and making decisions associated with the design and management of safety-critical systems and high-value assets. In the context of PRA, risk is characterised by the likelihood of an adverse event occurring and the severity of the consequences of the event. The likelihood of occurrence for uncertain adverse events is quantified through probabilistic event sequence and system modelling. The consequences and expected costs/gains are compared and evaluated by finding an appropriate utility metric – obvious examples include financial cost and loss of human life; however, in many applications these are overly simplistic [13]. Probabilistic risk assessment is applied in a range of industries including nuclear, aerospace and chemical process. Whilst the exact methodology used for conducting PRA differs between industries, they generally adhere to the key steps as outlined by the US Nuclear Regulatory Commission (USNRC) and the International Atomic Energy Agency (IAEA) [2]:

1. Initial information collection.
2. Event-tree development.
3. System modelling.
4. Reliability modelling.
5. Failure sequence quantification.
6. Consequence analysis.

Further detail is provided for each step in the following subsections.

3.1. Initial information collection

Information regarding the design and operation of the structure in question is collated. Details such as component specifications, loading and environmental conditions are considered. The information gathered at this stage is used to inform the subsequent steps. Given the large quantity of information required for conducting PRA, an important factor to be considered at this stage is the method by which the necessary information is represented, stored and managed. A common practice is to utilise a database [13]. There is a clear analogy here with the operational evaluation stage for SHM.

3.2. Event-tree development

Event trees outline potential accident sequences – combinations of initiating events and the subsequent system failures or successes that may result in an adverse consequence. The sequences of system failures and successes are known as *top events*. The system failures identified in the event-tree development stage are subsequently modelled as fault trees. An example event tree for a system designed to prevent injury following a jump from a plane is shown in Fig. 1.

3.3. Fault tree development

Fault trees are used in PRA to facilitate the quantification of system failure probabilities. The development of fault trees involves expressing the causal relationships between component failures and subsystem failures using Boolean logic gates. The level of detail captured in the fault tree (the level of components which are incorporated) is determined by the component level for which meaningful reliability data can be obtained. Failures of components belonging to the most fundamental level incorporated in the fault tree are known as *basic events* and are represented in a fault tree diagram as circles. Intermediate and top events are defined as combinations of other intermediate, and basic events through Boolean logic gates such as the AND-gate and OR-gate. The fault tree diagram notation for the AND-gate and OR-gate are shown in Fig. 2. An example fault tree for the deployment of a reserve parachute is shown in Fig. 3.

3.4. Reliability modelling

Information regarding the reliability of system components and the frequency of initiating events is necessary to quantify the probabilities of top events; this is typically gleaned from data and by applying appropriate reliability models.

3.5. Failure sequence quantification

By assigned the components in the fault trees with failure rates, the probability of top events may be computed. Propagating the probabilities of the initiating events and top events through the event tree allows the probability of each possible outcome in the event tree to be calculated.

3.6. Consequence analysis

With the probability of an adverse event quantified, a cost/utility metric must be chosen so that the risk associated with the failure sequence may be assessed. The risk assessment may then be used to inform design decisions, such as increasing safety by introducing additional redundancies in the system, or optimising cost by removing components that do not cause the risk to fall below an acceptable threshold. The risk assessment may also be used to inform risk-based inspection for a system that is in operation.

4. Probabilistic Graphical Models

Probabilistic graphical models (PGMs) are a powerful framework for reasoning and decision-making under uncertainty - a core problem in SHM. Probabilistic graphical models are representations of joint probability distributions, in which nodes

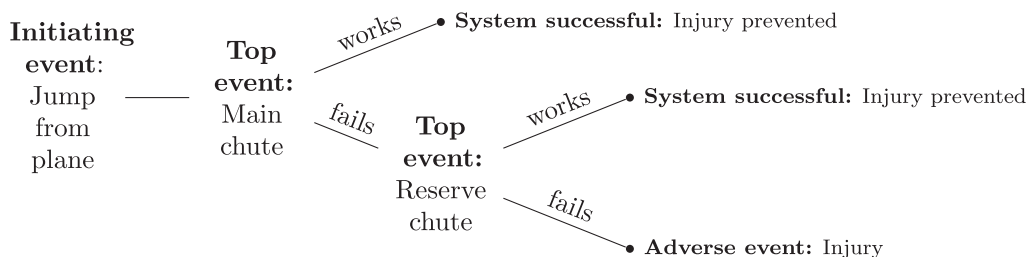


Fig. 1. An example event tree for a parachute system to prevent fall injuries [14].



Fig. 2. Fault tree diagram representations of the AND-gate (left) and OR-gate (right).

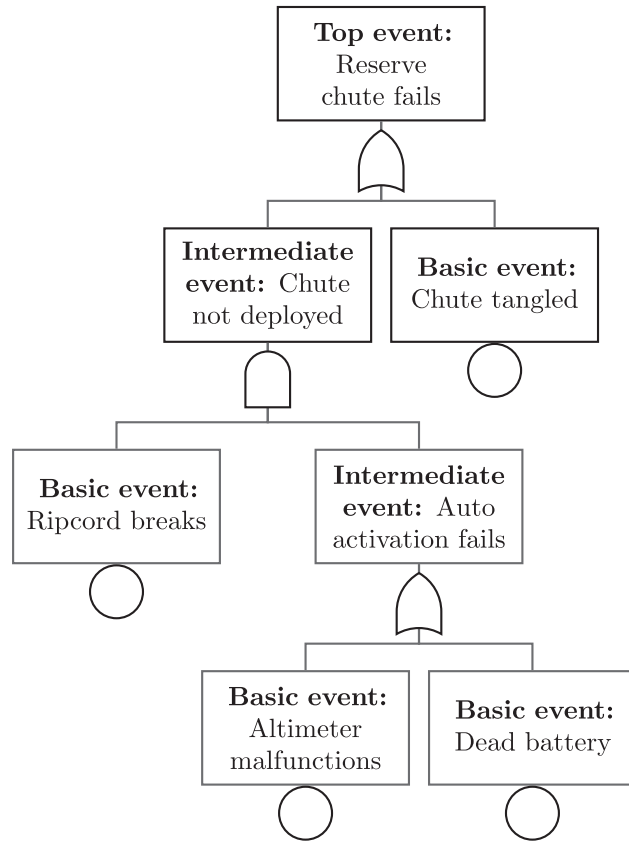


Fig. 3. An example fault tree for a reserve parachute system [14].

denote a set of random variables and edges connecting nodes imply dependency between variables. The probabilistic graphical model representation provides benefits over a flat (non-graphical) representation [15]:

- They provide a compact and intuitive representation of complex probability distributions which makes them easier to understand, communicate and learn.
- They facilitate efficient computation by exploiting local independence structures.

A PGM over a set of N variables X may be specified by a set of M local functions $f(Y_i)$, where Y_i is some subset of X , and a graph G comprised of nodes/vertices V and edges E . The joint probability distribution represented by the graph is obtained by:

$$P(X_1, X_2, \dots, X_N) = K \prod_{i=1}^M f(Y_i) \tag{1}$$

where K is a normalisation factor ensuring the probabilities sum to unity.

There are two classes of problem associated with PGMs: *inference* and *learning*. Inference is concerned with obtaining the marginal or conditional probabilities of a subset of variables Z given any other subset Y . i.e. $P(Z|Y)$. Learning is concerned with obtaining the graph structure and parameters given a complete, or incomplete, set of observed data values for X . i.e. $G, f(Y_i)|X$. The remainder of the current paper will be primarily concerned with inference problems and their application in a risk-informed SHM framework.

4.1. Bayesian networks

Bayesian networks (BNs) are a form of PGM. Specifically, they are directed acyclic graphs (DAGs) in which nodes represent random variables and edges connecting nodes represent conditional dependencies between variables. For discrete random variables, the local functions that describe the conditional probability distributions (CPDs) between variables are

conditional probability tables (CPTs), and in the case of continuous random variables are conditional probability density functions (CPDFs).

Fig. 4 shows a simple Bayesian network where X is a parent of Y and an ancestor of Z ; Z is said to be the child of Y and a descendant of X . Node X is independent of other nodes and so is specified by the unconditional distribution $P(X)$. Observed variables are shaded grey.

Given observations on a subset of nodes in a BN, inference algorithms can be applied to obtain posterior distributions over the unobserved random variables. In some cases, analytical solutions of posterior distributions may be found by using exact inference methods. To solve the inference problem using direct computation from the joint probability distribution, the computational complexity increases exponentially with the size of the graph and quickly becomes intractable. Fortunately, algorithms have been developed that allow efficient computation [16].

4.2. Influence diagrams

Bayesian networks can be augmented to represent decision processes by incorporating nodes for decision variables and utility functions – these augmented networks are known as *influence diagrams* [17]. Decision nodes are denoted by squares and utility nodes are denoted by rhombi, as shown in Fig. 5. Edges connecting random variables to utility nodes denote that the utility function is dependent on the state of that variable. Similarly, edges connecting decision nodes to utility nodes denote that the utility function has a dependence on the action decided upon. Edges connecting random variables or decision nodes to other decision nodes denote order, that is to say the random variable or decision is observed prior to the decision being made; such edges are referred to as *informational links* as they do not imply a functional dependence but rather that the information regarding the state of the variable is required for the decision to be made.

The influence diagram shown in Fig. 5 can be interpreted as a binary decision process regarding whether to go out for a walk or stay in and watch TV under uncertainty in the future weather condition W_c given an observed weather forecast W_f . The nodes W_f and W_c can be considered as binary random variables representing the weather forecast and actual weather condition, respectively, with possible states $domain(W_f) = domain(W_c) = \{bad, good\}$ and the weather forecast is dependent on the weather condition. The possible actions can be summarised as $domain(D) = \{TV, walk\}$. The utility U achieved is then dependent on both the weather condition experienced and the decided action.

In general, a policy δ is a mapping from all possible observations to possible actions. The problem of inference in influence diagrams is to determine an optimal strategy $\Delta^* = \{\delta_1^*, \dots, \delta_n^*\}$ given a set of observations on random variables where δ_i^* is the policy for the i^{th} decision to be made in a strategy Δ^* that yields the *maximum expected utility (MEU)*. The expected utility is a function of probability and utility; and by this definition is equivalent to risk.

5. Definitions

To establish a framework for mapping PRA onto SHM, some fundamental concepts will first be defined. In addition, a notation will be established for describing structures that can be expressed as hierarchical graphs.

One begins with the full structure of interest \mathbf{S} . It is assumed that \mathbf{S} may be decomposed into a discrete number of constituent elements, which are referred to as substructures. Substructures are considered to be entities which may, in principle, be available for independent testing prior to incorporation into the full-scale structure. This allows a hierarchical framework, where substructures may be further decomposed. At the base of the hierarchy are components, which are substructures which cannot be decomposed further. A joint is the physical mechanism by which substructures are joined together, thus moving up the hierarchy until \mathbf{S} itself is finally obtained. Joints themselves may be physically complex to the extent that they may be deemed substructures.

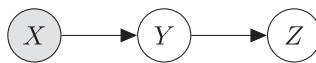


Fig. 4. An example Bayesian network.

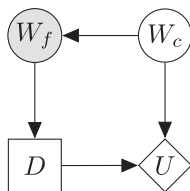


Fig. 5. An example influence diagram representing the decision of whether to go outside or stay in under uncertainty in the future weather condition given an observed forecast.

Fig. 6 depicts a graphical representation of a hierarchical structure that may be considered without loss of generality. Nodes represent the global structure and its constituent units and edges represent the dependence of a (sub) structure on its constituent units. At the top, or level 1, of the hierarchy is the global structure with the hierarchy level denoted in the superscript. It can be seen that the global structure S^1 is comprised of two substructures s_1^2, s_2^2 and a joint j_1^2 , i.e. $S^1 = \{s_1^2, j_1^2, s_2^2\}$. These units form the second level of the hierarchy. s_1^2 and s_2^2 may in turn be expanded to yield $S^1 = \{s_1^3, j_1^3, c_1^3, j_1^2, \{c_1^3, j_2^3, s_2^3\}\}$. Progressing down the hierarchy levels, one can continue to expand the substructures into constituent units until the L^{th} level of the hierarchy which is comprised solely of base units (components and joints). By taking the expansion of S^1 into its constituent base units and discarding the repeated units arising from substructures that share components, one obtains a list of the base units that form a given structure S . Within each level of the hierarchy, units are numbered via a subscript from 1 to N_u^i where N_u^i is the number of a constituent unit type u in the i^{th} level of the hierarchy. The notation u_n^i where i is an integer from 1 to L and n is an integer from 1 to N_u^i provides a unique identifier for each unit within a structure. To further elucidate this notation; with reference to Fig. 6, one can consider the joint units ($u = j$) within the third level of the hierarchy ($i = 3$) of which there are in total two ($N_j^3 = 2$). Thus, if one is required to reference the second joint in the third level of the structure's hierarchy, the notation j_2^3 may be used.

It is assumed that there exists a set of features v , observable from S , that are produced according to a generative latent state model with latent state $H^*(t)$, where $H^*(t)$ is the true health state of S and may be expressed in terms of the true health states of the constituent components and joints $h^*c_n^i(t)$ and $h^*j_n^i(t)$, respectively, i.e. $H^*(t) = \{h^*c_1^1(t), \dots, h^*c_{N_c^L}^L(t), h^*j_1^2(t), \dots, h^*j_{N_j^L}^L(t)\}$.

The structure S also has a predicted time-dependent health state vector $H(t) = \{hc_1^2(t), \dots, hc_{N_c^L}^L(t), hj_1^2(t), \dots, hj_{N_j^L}^L(t)\}$.

Health-state vectors can be constructed from any subset of components and joints.

For the structure/system S , there must exist a set of failure modes of interest $F = \{F_1, \dots, F_{N_f}\}$ whereby S ceases to be fit for purpose. It is assumed that a given failure mode is dependent on the health states of a subset of components, joints and substructures for which a health-state vector can be constructed. In addition, each failure state has an associated utility U_{F_n} .

Finally, for the structure S , there also exists a set of decisions $d = \{d_1, \dots, d_{N_d}\}$ which affect H^* , each having an associated utility U_{d_i} . In addition, there will exist some set of environmental conditions $e = \{e_1(t), \dots, e_{n_e}(t)\}$ that will alter the distribution of v .

6. Mapping PRA onto SHM

Upon examination, it becomes apparent that there are both differences and similarities between the paradigms for SHM and PRA that can be examined to determine which aspects of PRA will be useful for SHM. Whilst it is clear that both SHM and PRA are utilised for the purpose of making decisions in the face of uncertainty, PRA is conducted offline for a system experiencing a set of anticipated initiating events. In contrast, the decision processes for which SHM is implemented are online and require continual predictions of the damage state of the structure. It is for this reason that the event-tree development stages and failure sequence quantification stages in PRA are less applicable to SHM.

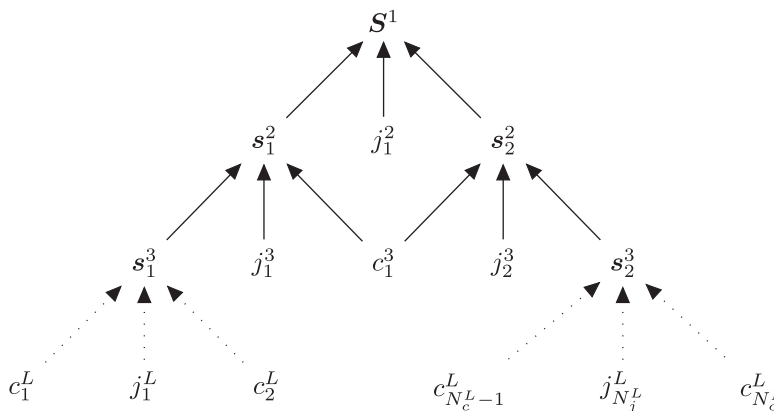


Fig. 6. A hierarchical graphical representation of a generic structure S . Superscripts denote the level in the hierarchy and subscript indexes each type of constituent unit in a given level. Dotted edges imply an arbitrary structuring between levels.

Both paradigms begin with collating information regarding the structure and defining the context in which decisions are to be made. In fact, the first three stages of the PRA paradigm involve expressing the structure and context in a logical way which facilitates the quantification of risk and the decision-making process. It is in this formal expression of the structure that the decision-making process in the SHM paradigm stands to benefit. An overview of the risk-based SHM paradigm is as follows:

1. Operational evaluation.
2. Failure-mode modelling.
3. Decision modelling.
4. Data acquisition.
5. Feature selection.
6. Statistical modelling for feature discrimination.

6.1. Operational evaluation

With the aim of justifying the use and defining the context of a risk-based SHM system, the operational evaluation stage seeks to answer many of the same questions as in the standard paradigm. However, some questions require an approach that facilitates the failure-mode modelling and decision-modelling stages.

Foremost, information regarding the components \mathbf{c} , joints \mathbf{j} , substructures \mathbf{s} , and the dependencies between them is required.

When identifying the critical damage states of the structure \mathbf{S} , one should aim to identify the failure modes of interest \mathbf{F} . Critical components, joints and substructures/subsystems that contribute to \mathbf{F} should also be identified at this stage. The predicted damage states of these components h should be defined. The damage states of the critical substructures/subsystems \mathbf{H} should be defined as a vector in terms of h .

For each failure mode in \mathbf{F} , potential decisions \mathbf{d} should be identified and the ways in which the actions influence the structure or likelihood of failure modes occurring should be determined. Utility values \mathbf{U}_F and \mathbf{U}_d for all \mathbf{F} and all \mathbf{d} , respectively, should be quantified. The selection of utility values will determine the behaviour of the decision-making agent, and is analogous to setting a decision threshold in a standard SHM paradigm.

Environmental influences \mathbf{e} should also be identified. It should also be decided whether the SHM system is to evaluate the health of the structure at static, independent instances in time, or predict future health states, thereby requiring a model forecasting the degradation of the structure.

For large, complex structures it may be beneficial to borrow the data management techniques used in PRA, such as databases, to organise the information obtained during the Operational Evaluation stage. This will allow for a rigorous and structured approach to the information collection and allow for the identification of aspects of the SHM system that require further specification or more information. Having a formal information structure will also expedite the subsequent failure-mode modelling step which requires detailed knowledge of the physical structure.

6.2. Failure-mode modelling

For each of the failure modes of interest in \mathbf{F} , one should proceed to construct a fault tree, such as that shown in Fig. 7, based upon the health states of the relevant components, joints and substructures/subsystems. It should be noted here that, in many cases, the exact nature of the failure modes will be unknown and so a best estimate based on engineering judgement may be used.

Fault trees offer a rigorous and consistent structure for expressing the failure modes; however, as statements in Boolean logic they are limited in their flexibility. In the context of SHM, it is desirable to represent the components in a fault tree as having multiple damage states, and it is for this reason that one should map the constructed fault trees into Bayesian networks. Bobbio et al. outline a convenient mapping from fault trees into Bayesian networks in [18], whilst also highlighting the additional flexibility that is granted by doing so. Additionally, Bayesian networks are used to represent structural failures in [19].

In the example shown in Fig. 8, the component health states, substructure health states and failure event are represented as random variables where the substructure health states are conditioned on the component health states and the failure events is conditioned on the substructure health state. The random variables are defined using a conditional probability distribution (CPD) which may be discrete or continuous.

A node representing the health-state vector of the critical components and joints \mathbf{H} should be included in the fault tree, as this latent state will be predicted during the statistical modelling process. To define the vector \mathbf{H} within the Bayesian network, the conditional dependence between the nodes representing the local health states of the components and joints and \mathbf{H} are expressed as a binary logic table.

One function of the failure-mode Bayesian network is to allow the flow of information from the statistical model to the decision, whilst parsing the information in a way that facilitates the defining of the failure events \mathbf{F} . The network also allows

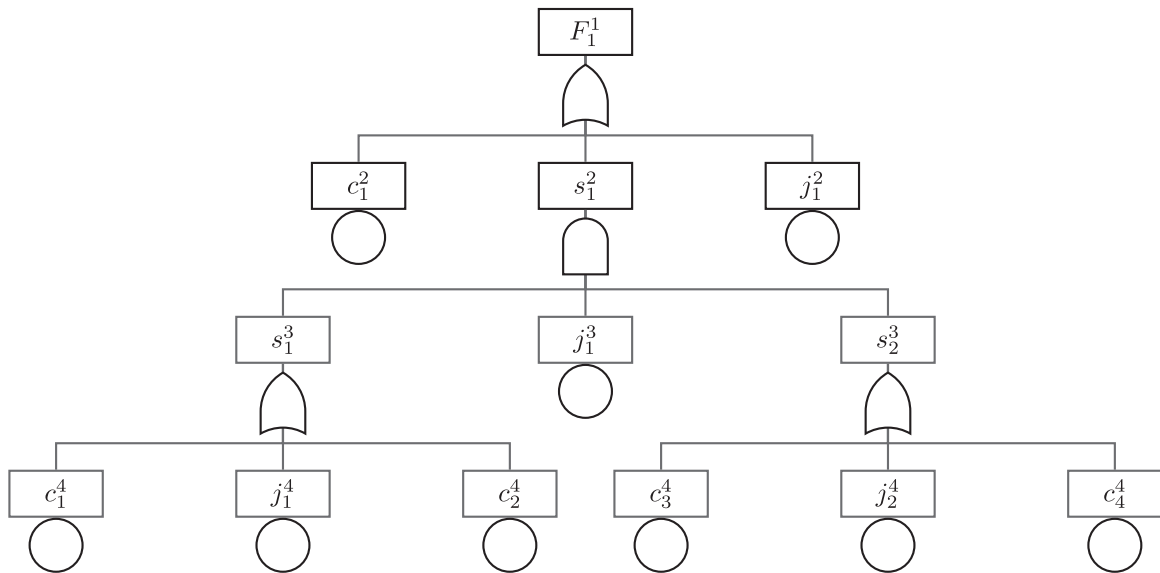


Fig. 7. A fault tree of a (single) failure mode F_1 where the superscript denotes the hierarchy level and the subscript is an identifier.

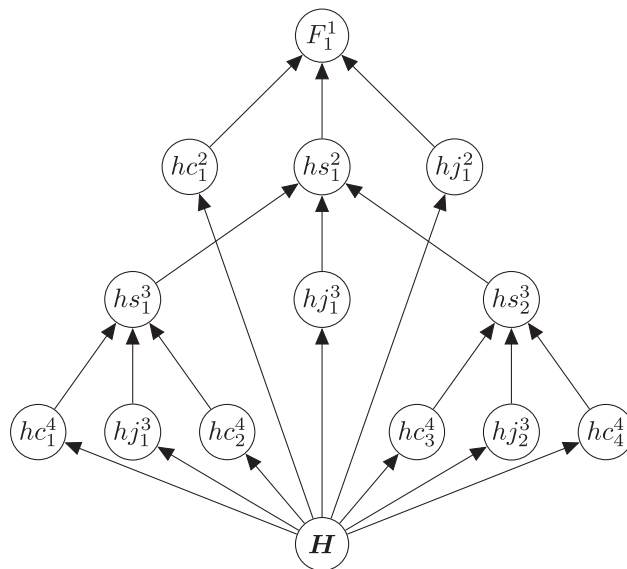


Fig. 8. A Bayesian network of failure mode in F_1 .

the computation of marginal distributions for the probability of failure in each component, joint, or substructure allowing for damage localisation.

6.3. Decision modelling

Modelling the decision process involves augmenting the Bayesian network developed in the failure modelling stage with nodes for each decision in \mathbf{d} and for utilities \mathbf{U}_F and \mathbf{U}_d to produce an influence diagram. Decision nodes in which the actions alter the probability of a state or event should modify the CPDs accordingly. Utility nodes are constrained to be leaf nodes and should be dependent on the appropriate failure events or decisions.

For static problems, it may be convenient to model the decision process in a separate influence diagram which receives information regarding failure probabilities from the fault tree. This issue is because it is implicit that the decision is made after observations are made; if one attempts to solve a network in which a decision is made that yields a state that is inconsistent with the observed state, a conflict arises.

6.4. Data acquisition

The data acquisition process should not differ from that in the standard SHM paradigm. Here, there is a subtlety that the data acquisition system should be designed so as to optimise the decision-making rather than damage identification.

6.5. Feature selection

The feature selection process should not differ from that in the standard SHM paradigm. Again, there is the subtlety that the features should be selected so as to optimise the decision-making.

6.6. Statistical modelling

The purpose of the statistical model is to predict the critical health states \mathbf{H} given the selected feature set \mathbf{v} . As aforementioned, it is assumed that \mathbf{v} is produced through a generative latent state model, with latent-state \mathbf{H} . Probabilistic classifiers that output a probability distribution over all possible states of \mathbf{H} , such as Gaussian mixture models (GMMs), are compatible. Probabilistic classifiers are instrumental in building robustness to the uncertainty surrounding the true health state of \mathbf{S} into the decision process. Ideally, the chosen statistical model will be capable of consistently identifying the actual health state under all identified operating and environmental conditions \mathbf{e} , or at least appropriately reflect the uncertainty caused by varying conditions in the prediction.

Finally, if a model describing the degradation of \mathbf{S} (i.e. a transition model for \mathbf{H}) is required for forecasting failure events in the time-dependent case, the CPDs defining $P(\mathbf{H}_t | \mathbf{H}_{t-1}, \mathbf{d})$ should be specified accordingly.

7. Case study: Four-bay truss

To demonstrate the probabilistic graphical model formulation of a risk-based approach to an SHM problem, it was applied to a four-bay truss identical to that used in [20] and shown in Fig. 9. For clarity, the example will be limited to a single failure mode and a single binary decision. The failure of joints will also be ignored.

The truss was composed of 20 aluminium members each with length 250 mm and cross-sectional area 177 mm²; the overall length of the structure being 1 m and the height being 0.25 m. The members were pinned together with steel bolts in lubricated holes. The truss was subjected to a preload of 5 kg at point P and additional consecutive loads of 10, 20 and 30 kg at each of the 8 points L in turn. For each of the 24 load cases, microstrains were measured at the midpoints of the 12 horizontal and vertical members. This process was performed 8 times in total, once with each cross-member removed.

In addition to the experimental data, a finite element simulation of the truss was developed whereby removal of a cross-member was achieved by assigning a Young’s modulus of 1 MPa. As well as the 10, 20 and 30 kg loads used in the experiment, the truss was simulated with loads of 5, 15 and 25 kg. Furthermore, strains were obtained for the truss under each load case in its undamaged condition, i.e. with all cross-members intact.

7.1. ‘Operational Evaluation’ of Truss

In order to construct a risk-based decision framework for the truss, one must first define it formally. As it was elected to ignore joints, the global truss structure \mathbf{T} can be defined as four substructures, one for each bay i.e. $\mathbf{T} = \{\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3, \mathbf{b}_4\}$. As only the failure of cross-members was considered, each bay can in turn be defined as two components, for example, $\mathbf{b}_1 = \{m_9, m_{13}\}$. Consequently, the global structure may be represented as $\mathbf{T} = \{m_9, m_{10}, m_{11}, m_{12}, m_{13}, m_{14}, m_{15}, m_{16}\}$.

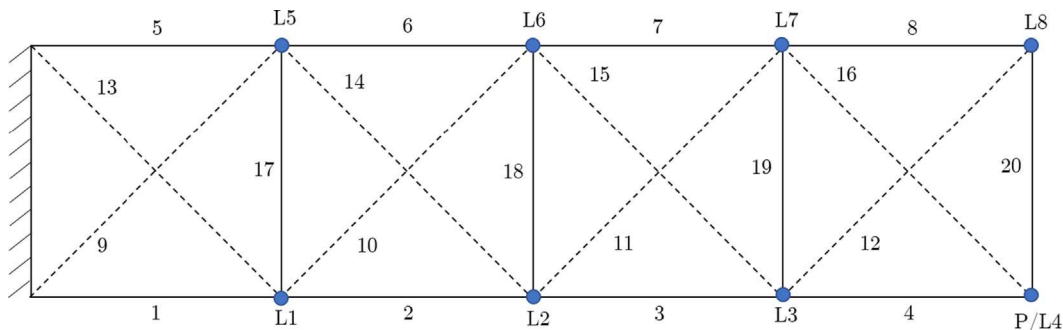


Fig. 9. A 2-dimensional four-bay truss comprised of 20 members, 8 of which are removable and denoted by a dashed line. Loads are applied at points L, and a preload is applied at point P. Load positions are shown as blue dots. The bays are numbered left to right from 1 to 4.

A single failure mode F_T of the truss was considered; the full or partial collapse of the structure. This failure mode corresponds to the event where the truss is no longer able to support the load/preload, hence, F_T occurs when both cross-members in a single bay fail.

In an attempt to minimise the occurrence of the failure mode of interest F_T , a single binary decision d was identified; a choice between ‘do nothing’ and ‘perform maintenance’. In addition, utilities were assigned to the failure event and the decidable actions somewhat arbitrarily, though, in a manner which intends to reflect the relative costs associated with failure and maintenance in real-world engineering applications. There is a small positive utility associated with the structure being in its operational state, and a large cost associated with the structure being in its failed state. The ‘do nothing’ action is given zero utility and the ‘perform maintenance’ action is given a cost greater than utility gained through the structure being operational in a given time-step, although less than the cost of failure. The utilities assigned to the failure and decision are shown in Tables 1 and 2, respectively. Further discussion of the challenges associated with assigning utilities is provided in Section 8.2.

For the purposes of demonstration it is assumed that the load on the truss will be uncertain, varying in discrete time within the interval $[0, w_{max}]$ where w_{max} is defined in subSection 7.3. Furthermore, it is assumed that the location of the load also varies in discrete time and that, in the limit of infinite time-slices, each of the 8 locations is visited an equal number of times.

7.2. Truss failure modelling

The failure mode F_T of the truss can be represented as the fault tree shown in Fig. 10 where the failure of a bay \mathbf{b} is defined as the and-gate of two cross-member m failures, and the failure of the truss F_T is defined as the or-gate of the bay failures.

To map the fault tree for the failure event F_T into a probabilistic graphical model, \mathbf{hb} and hm will be used to denote the random variables that represent the local binary health states of the bays and cross-members, respectively, where 0 corresponds to intact and 1 corresponds to failed. Additionally, \mathbf{H} will be used to denote the random variable vector for the health state of the global structure where $\mathbf{H} = \{hm_9, hm_{10}, hm_{11}, hm_{12}, hm_{13}, hm_{14}, hm_{15}, hm_{16}\}$. For conciseness, the vector \mathbf{H} will, on occasion, be summarised as $\mathbf{H} = H$, where H is the decimal representation of the 8-bit binary number specified by the vector. Finally, the F_T notation will be retained to represent the random variable corresponding to the failure event. The Bayesian network corresponding to the fault tree shown in Fig. 10 is shown in Fig. 11. The conditional probability distributions specifying $P(F_T|\mathbf{hb}_1, \mathbf{hb}_2, \mathbf{hb}_3, \mathbf{hb}_4)$ (or $P(F_T|\mathbf{hb})$ for brevity) and $P(\mathbf{hb}_1|hm_9, hm_{13})$ (or $P(\mathbf{hb}_1|\mathbf{hm}_{b_1})$ for brevity) are shown in Tables 3 and 4, respectively.

7.3. Health-state transition modelling

The purpose of health-state transition modelling is to develop the conditional probability distribution $P(\mathbf{H}_{t+1}|\mathbf{H}_t, d)$ that predicts the future health state of the truss forward in time, given the current health state and the decided action.

For the purpose of this demonstration, it was decided that the ‘perform maintenance’ action simply returns the structure to its undamaged state, i.e. with no cross-members failed, with probability 1, independent of \mathbf{H}_t .

With regard to the ‘do nothing’ action, it was first assumed that the truss would not spontaneously transition from a more advanced damaged state to a lesser one, that is to say, cross-members would not self-repair in the absence of intervening maintenance, or, without maintenance the health-state of the structure monotonically degrades as a function of time.

The assumed loading range $[0, w_{max}]$ was discretised into 100 evenly-spaced increments; combined with the 8 possible load locations, this resulted in 800 unique considered load cases L_c .

For a given load case L_c , a transition in health state was defined as $\mathbf{H}_{t+1} = \mathbf{H}_t + \delta\mathbf{H}_{t \rightarrow t+1}$ where $\delta\mathbf{H}_{t \rightarrow t+1}$ is an 8-bit binary vector with i^{th} entry equal to 1 if the yield stress of aluminium (300 MPa) is exceeded in member m_{i+8} when the truss is simulated in health state H_t subject to load case L_c , and equal to 0 otherwise. The conditional probability of transitioning from $\mathbf{H}_t = H_t$ to $\mathbf{H}_{t+1} = H_{t+1}$, $P(\mathbf{H}_{t+1} = H_{t+1}|\mathbf{H}_t = H_t, L_c)$, was assigned unity if $\delta\mathbf{H}_{t \rightarrow t+1} = \mathbf{H}_{t+1} - \mathbf{H}_t$ for $\mathbf{H}_t = H_t$ and $\mathbf{H}_{t+1} = H_{t+1}$ for load case L_c , and assigned zero otherwise. The full transition matrix $P(\mathbf{H}_{t+1}|\mathbf{H}_t)$ was then populated where the entry $P(\mathbf{H}_{t+1} = H_{t+1}|\mathbf{H}_t = H_t)$ is given by,

$$P(\mathbf{H}_{t+1} = H_{t+1}|\mathbf{H}_t = H_t) = \frac{\sum_{L_c=1}^{N_{L_c}} P(\mathbf{H}_{t+1} = H_{t+1}|\mathbf{H}_t = H_t, L_c)}{N_{L_c}} \tag{2}$$

where N_{L_c} is the total number of load cases considered and $N_{L_c} = 800$.

Table 1

A table showing the entries of the utility function $U(F_T)$ where $F_T = 0$ and $F_T = 1$ denote the truss being operational and failed, respectively.

F_T	$U(F_T)$
0	15
1	-285

Table 2

A table showing the entries of the utility function $U(d)$ where $d = 0$ and $d = 1$ denote the 'do nothing' and 'perform maintenance' actions, respectively.

d	$U(d)$
0	0
1	-100

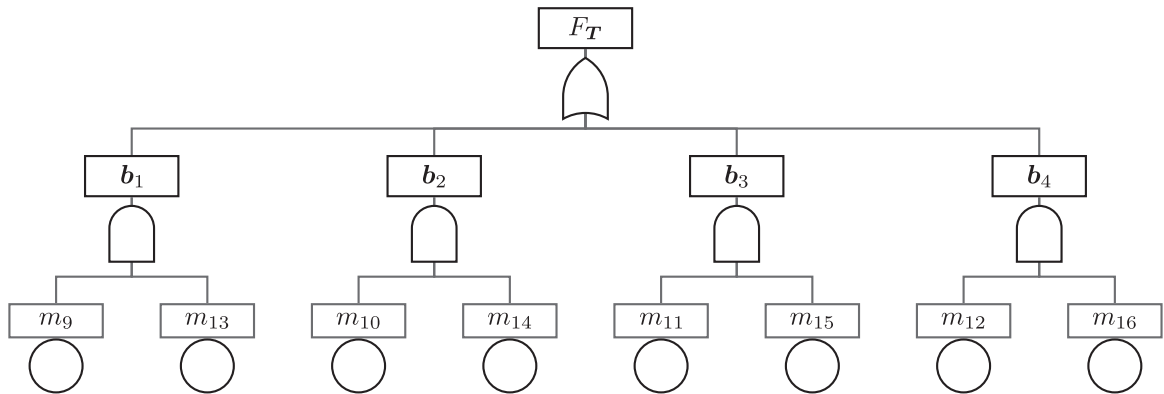


Fig. 10. A fault tree of failure mode F_T for a four-bay truss. The failure mode F_T occurs if at least one bay b fails. A bay b will fail if both cross-members m fail.

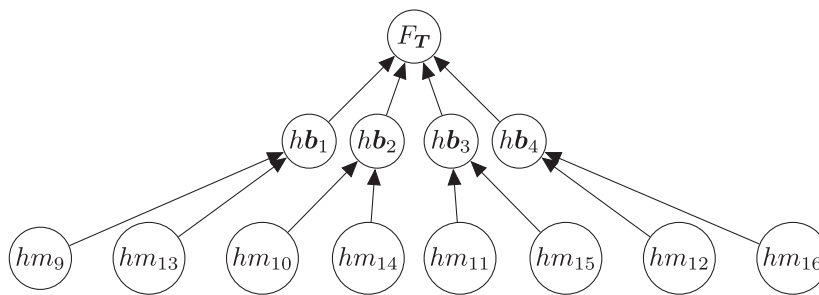


Fig. 11. A Bayesian network representation of failure mode F_T for a four-bay truss.

Table 3

A table showing the entries of the conditional probability distribution $P(F_T|hb_1, hb_2, hb_3, hb_4)$ where $hb = 0$ and $hb = 1$ denote a bay being intact and failed, respectively, and $F_T = 0$ and $F_T = 1$ denote the truss being operational and failed, respectively.

hb_1	hb_2	hb_3	hb_4	$P(F_T = 1 hb)$	$P(F_T = 0 hb)$
0	0	0	0	0	1
1	0	0	0	1	0
0	1	0	0	1	0
1	1	0	0	1	0
0	0	1	0	1	0
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
0	1	1	1	1	0
1	1	1	1	1	0

Table 4

A table showing the entries of the conditional probability distribution $P(hb_1|hm_9, hm_{13})$ where $hm = 0$ and $hm = 1$ denote a member being intact and failed, respectively, and $hb = 0$ and $hb = 1$ denote a bay being intact and failed, respectively.

hm_9	hm_{13}	$P(hb_1 = 1 hm_{b_1})$	$P(hb_1 = 0 hm_{b_1})$
0	0	0	1
1	0	0	1
0	1	0	1
1	1	1	0

For illustrative purposes, the maximum load w_{max} was determined by asserting $P(\mathbf{H}_{t+1} \neq 0 | \mathbf{H}_t = 0) = 0.005$ and the value of w_{max} that satisfied the condition was found to be approximately 6900 kg. This is obviously somewhat arbitrary, and in practice the maximum load for a structure may be estimated during the operational evaluation stage.

The transition model developed provides a means of forecasting future health states of the truss, a heatmap representation of the transition matrix is shown in Fig. 12.

7.4. Statistical classifier development and training

The purpose of the statistical classifier is to obtain a probability distribution over the current health state given a set of observed features, i.e. $P(\mathbf{H}_t | \mathbf{v})$. In the current case study, the damage indicative features are the strains measured from the horizontal and vertical members of the truss. The classifier selected for the current case study was comprised of two components; a detector and a localiser, corresponding to the first two stages of Rytter's hierarchy. Whilst generative models may better reflect the causality of the problem at hand, where it is assumed that the features are generated as a result of the latent state and $P(\mathbf{H}_t | \mathbf{v})$ may be computed via Bayes' Theorem, discriminative classifiers that directly learn a mapping from the feature space to the label space are also applicable in the risk-based decision framework.

A Gaussian novelty detector was implemented to determine the probability that the structure is currently in its undamaged state $P(\mathbf{H}_t = 0 | \mathbf{v})$, and, as its complement, the probability that the structure is currently in a damaged state $P(\mathbf{H}_t \neq 0 | \mathbf{v})$. The first two principal components of the simulated strain data for the undamaged structure are compared to those from the damaged structure in Fig. 13. Inspection of Fig. 13 reveals that it is possible to discriminate between the undamaged and damaged finite element simulation data using only the first principal component of the strains hence the novelty detector uses the first principal component as the discriminative feature. As such this principal component projection was learned from the training data and the detector was formed by computing the mean μ and standard deviation σ of the univariate distribution of the first principal component. If the first principal component of an incoming set of strains were to lie within the range $\mu \pm 3\sigma$, it was asserted that the observed strains came from the structure in its undamaged condition with confidence 0.997, i.e. $P(\mathbf{H}_t = 0 | \mathbf{v}) = 0.997$. If the first principal component lay outside the $\pm 3\sigma$ confidence interval, $P(\mathbf{H}_t = 0 | \mathbf{v})$ was given by the probability mass in the tail of the Gaussian probability density function parametrised by μ and σ .

The function of the localiser component of the statistical classifier is to distribute $P(\mathbf{H}_t \neq 0 | \mathbf{v})$ over the remaining 255 health states corresponding to the various combinations of cross-member failures. For consistency with [20], the classifier selected for this purpose was an artificial neural network (ANN) with an input node for each of the twelve strain measurements, an output node for each cross-member, and three hidden layers with twelve, twelve and eight nodes, respectively. The activation function used was the hyperbolic tangent function. As an identical classifier to that used in [20] was implemented, only health states $\mathbf{H}_t = \{1, 2, 4, 8, 16, 32, 64, 128\}$ are considered by the localiser in this case study.

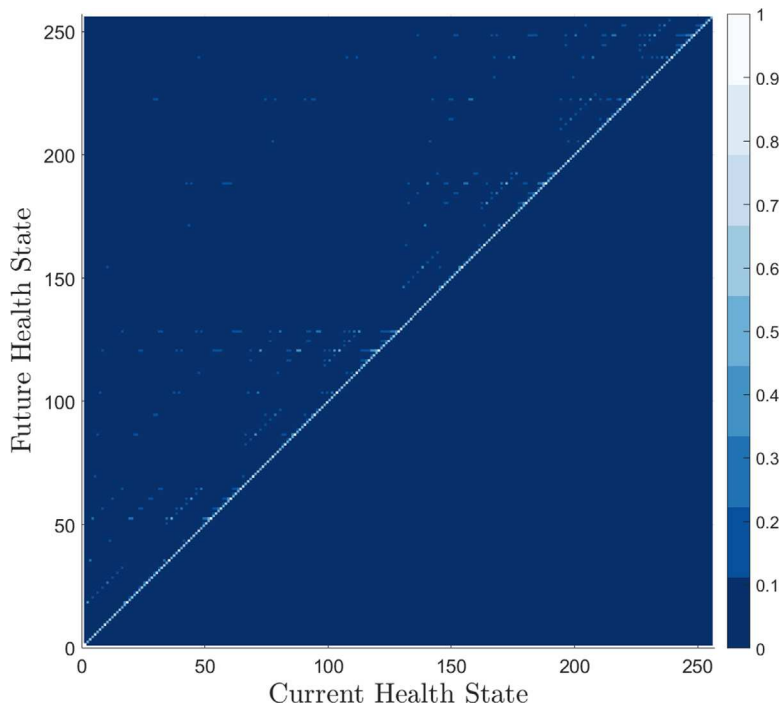


Fig. 12. A heatmap showing the conditional probability distribution transition matrix $P(\mathbf{H}_{t+1} | \mathbf{H}_t)$.

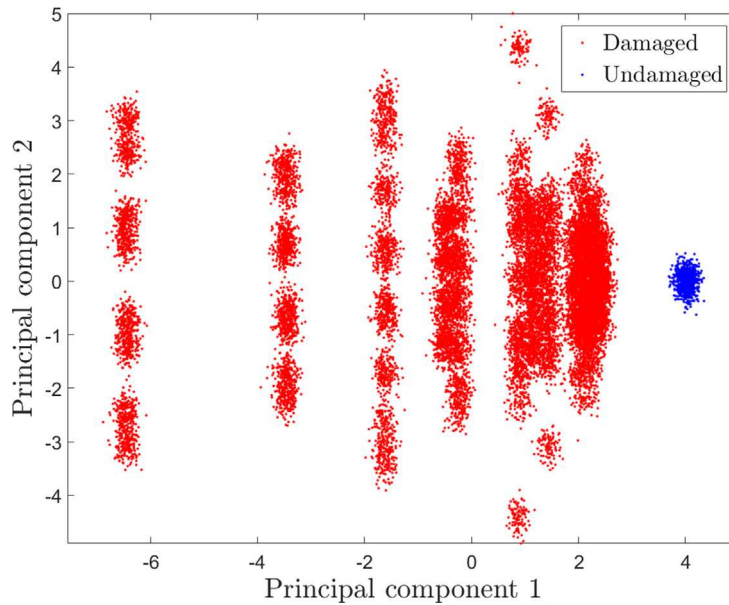


Fig. 13. A comparison of the distributions of the first two principal components of the strain data obtained from the finite element model of the undamaged and damaged truss.

In accordance with [20], a training dataset was constructed from the finite element data for the loads of 10, 20 and 30 kg and a validation dataset from the finite element data for the loads of 5, 15 and 25 kg. In both instances, 100 repetitions of the datasets were produced and superimposed with a noise pattern of 1 microstrain RMS. The optimal weights were computed using the scaled conjugate gradient (SCG) back-propagation algorithm [21] and evaluated and selected based on the classification performance of the network on the validation dataset. Although not inherently probabilistic, a pseudo-probabilistic interpretation for the activations of the output nodes was acquired through the use of a softmax function.

7.5. Overall probabilistic graphical model

The failure model, transition models, classifier, decisions and utilities can be combined to form a partially-observable Markov decision process represented by the limited memory influence diagram (LIMID) shown in Fig. 14. Fig. 14 shows the decision process for two decisions over three time-slices. The edge connecting \mathbf{v}_{t+0} to d_{t+0} implies that the features are observed prior to the decision being made. Similarly, the edge connecting d_{t+0} to d_{t+1} implies d_{t+0} is decided before d_{t+1} .

It should be noted that the model shown in Fig. 14 assumes a generative model for the features \mathbf{v} ; for discriminative classifiers the direction of the conditioning edge connecting nodes \mathbf{H}_{t+0} and \mathbf{v}_{t+0} would be reversed.

7.6. Results

7.6.1. Novelty detector

When applied to the experimental strain data from the damaged truss, the univariate Gaussian novelty detector was able to correctly identify 175 out of the 192 observations as novel with respect to the simulated undamaged data, thereby yielding an overall accuracy of 91.1% and shown in the confusion matrix in Fig. 15.

The 8.9% misclassification can be elucidated by examining the distribution of the experimental data as projected through the principal component mapping learned from the finite element simulation data, shown in Fig. 16. Separability between the projected damaged and undamaged datasets is absent where values of the first principal component are approximately 3.9.

7.6.2. Localiser

When applied to the experimental strain data, the neural network localiser has an overall classification accuracy of 59.9%. The full confusion matrix is shown in Fig. 17. Whilst 60% classification accuracy may be considered low, for a 8 class problem the neural network provides a significant improvement over simply guessing which would yield an accuracy of 12.5%. The imperfect classifier has been deliberately chosen here as a possible source of uncertainty.

The misclassification error of 40.1% can be explained by considering the physics of the problem at hand. The selected damage sensitive features were the strains in the horizontal and vertical members of the truss; however, significant changes in the strains are only expected of members in the load path between the end fixture and applied mass. Therefore, the features are largely insensitive to damage when the mass is closer to the fixture than the damaged cross-member and any dif-

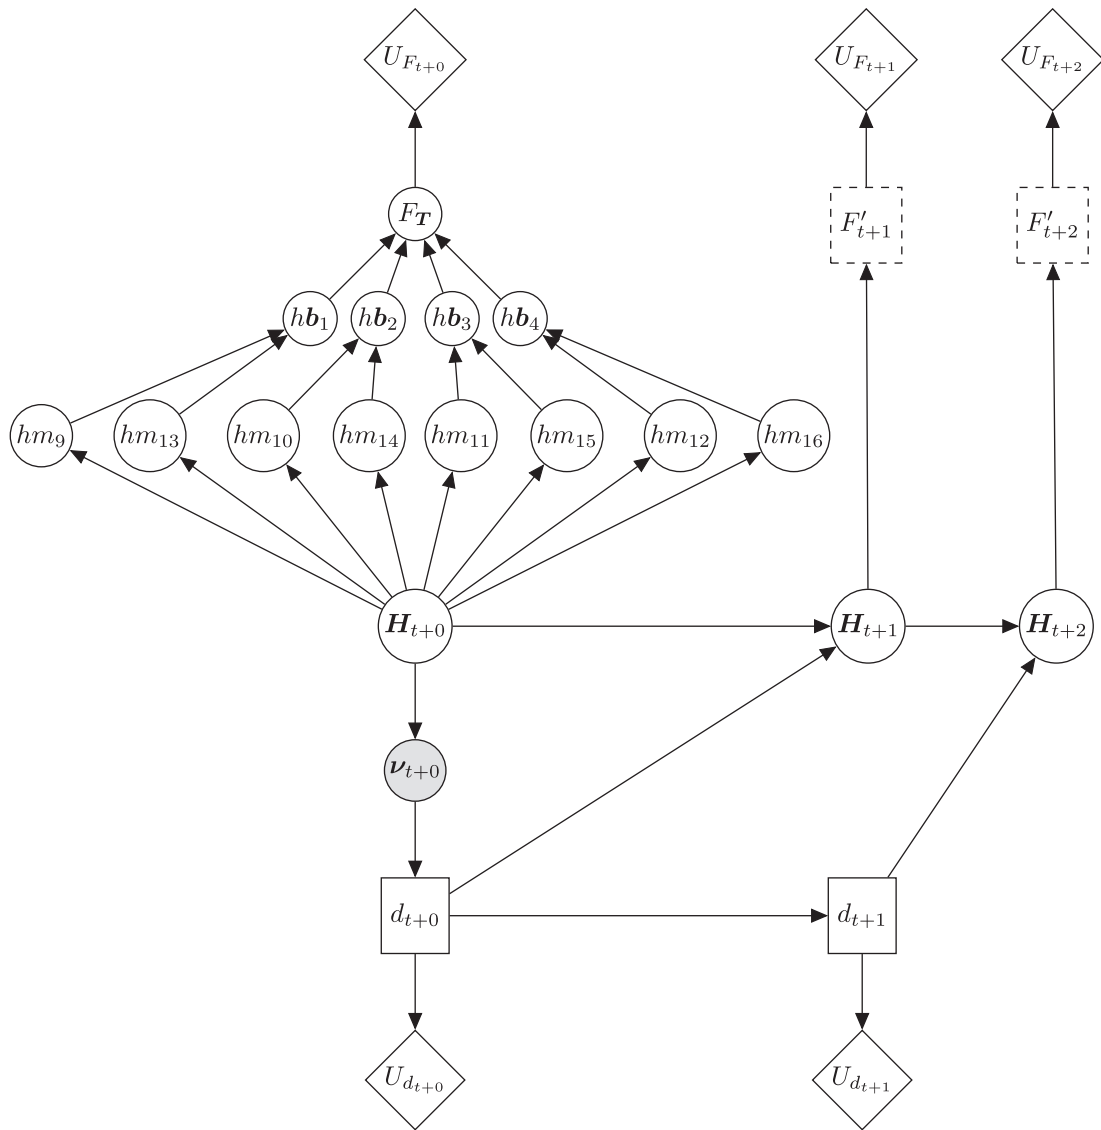


Fig. 14. An influence diagram representing the partially observable Markov decision process for determining the utility-optimal maintenance strategy for the cross-members of a four-bay truss given observations of strains made from the horizontal and vertical members. Observed variables are shaded grey. The fault tree failure models for the latter time steps have been represented as the nodes F'_i for compactness.

ferences can be attributed to the strains induced by the preload. In 72 of the 192 (37.5%) cases, the damaged cross-member is not in the load path.

7.6.3. Decision process results

The decision algorithm was tested on a dataset comprised of the experimental strains for the truss in its damaged conditions and, due to the lack of experimental data, finite element simulation data for the truss in its undamaged condition. Equal proportions of the undamaged and damaged data were used; 192 sets of strains from each.

The decision algorithm used for testing was similar to that shown in Fig. 14, except using the discriminative pseudo-probabilistic ANN classifier rather than a generative model. The graphical was modelled in MATLAB using the Bayes Net Toolbox [22] and solved using the junction tree algorithm for influence diagrams described in [23]. Utilities were as shown in Tables 1 and 2. Three failure events in consecutive time-steps were considered in conjunction with two 'do nothing'/'perform maintenance' decisions in the first two time-steps with a single observation made during the first time-step.

The performance of the decision process was evaluated with a metric similar to that of a classifier's overall accuracy. Whereas a classifier's accuracy is a comparison between the predicted outputs and the target outputs, this 'decision accuracy' is a comparison between the decided actions given the outputs of the classifier used to the optimal decided actions given perfect information of the health state, i.e. the utility-optimal decisions when the targets of the classifier are provided in

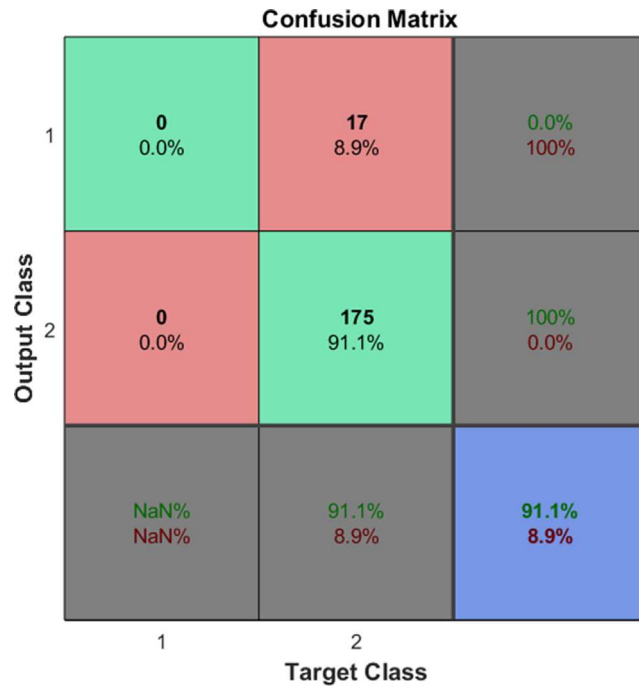


Fig. 15. A confusion matrix detailing the classification performance of the novelty detector on the experimental data.

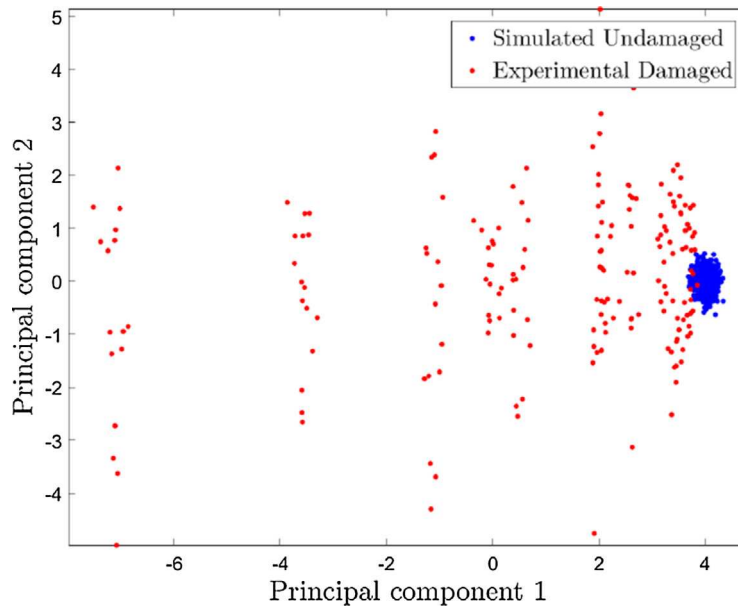


Fig. 16. A comparison of the distributions of the first two principal components of the strain data obtained from the finite element model of the undamaged truss and the strain data from the experiment performed on the damaged truss mapped through the projection learned from the simulated training data.

place of $P(H|y)$. The target and output classes ‘0’ and ‘1’ for Figs. 18–20 correspond to the ‘do nothing’ and ‘perform maintenance’ actions, respectively.

Fig. 18 shows the performance of the decision algorithm across all 786 decisions associated with the test dataset. It can be seen that an overall ‘accuracy’ of 93.2% was achieved meaning that the optimal decision given perfect information of the health state selected in 716 of the cases when the statistical classifier was used to infer the health state. In 40 of the cases, the ‘perform maintenance’ action was selected unnecessarily, this is a result of the uncertainty in the health state triggering a more, conservative action to be taken; this form of error is akin to a ‘false positive’ or type I error. In 12 of the cases the ‘do nothing’ action was deemed to be optimal whereas, had perfect information of the structures health state been available, the

Confusion Matrix

Output Class	1	23 12.0%	0 0.0%	2 1.0%	4 2.1%	0 0.0%	3 1.6%	4 2.1%	16 8.3%	44.2%
	2	0 0.0%	16 8.3%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	100%
	3	0 0.0%	0 0.0%	11 5.7%	0 0.0%	5 2.6%	0 0.0%	0 0.0%	0 0.0%	68.8%
	4	0 0.0%	0 0.0%	3 1.6%	15 7.8%	2 1.0%	0 0.0%	3 1.6%	3 1.6%	57.7%
	5	0 0.0%	0 0.0%	0 0.0%	0 0.0%	15 7.8%	0 0.0%	0 0.0%	0 0.0%	100%
	6	0 0.0%	2 1.0%	7 3.6%	4 2.1%	0 0.0%	21 10.9%	3 1.6%	3 1.6%	52.5%
	7	1 0.5%	4 2.1%	1 0.5%	1 0.5%	2 1.0%	0 0.0%	14 7.3%	2 1.0%	56.0%
	8	0 0.0%	2 1.0%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	0 0.0%	0.0%
			95.8% 4.2%	66.7% 33.3%	45.8% 54.2%	62.5% 37.5%	62.5% 37.5%	87.5% 12.5%	58.3% 41.7%	0.0% 100%
		1	2	3	4	5	6	7	8	
		Target Class								

Fig. 17. A confusion matrix detailing the classification performance of the artificial neural network localiser on the experimental data.

Confusion Matrix

Output Class	0	584 76.0%	12 1.6%	98.0% 2.0%
	1	40 5.2%	132 17.2%	76.7% 23.3%
		93.6% 6.4%	91.7% 8.3%	93.2% 6.8%
		0	1	
		Target Class		

Fig. 18. A confusion matrix detailing the 'accuracy' of the decision algorithm for both decisions in the three time-slice problem for $U(F_T = 1) = -285$.

optimal decision would, in fact, have been 'perform maintenance'. This form of error is akin to a 'false negative' or type II error.

The severity/significance of type I and type II errors is dependent on the context of the SHM system. For example, for an offshore wind structure, erroneously sending inspection/maintenance engineers has a higher cost relative to the failure event, whereas for a bridge the cost of inspection/maintenance is relatively lower with respect to the cost of failure.

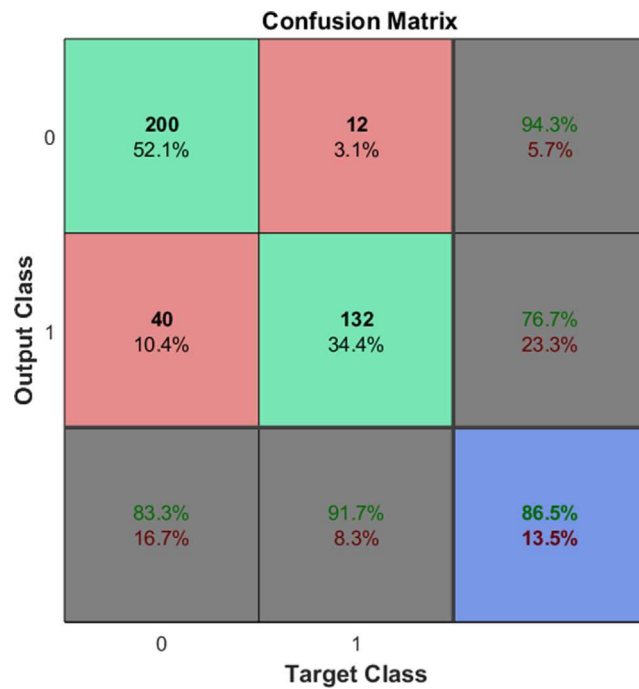


Fig. 19. A confusion matrix detailing the 'accuracy' of the decision algorithm for the first decision in the three time-slice problem for $U(F_r = 1) = -285$.

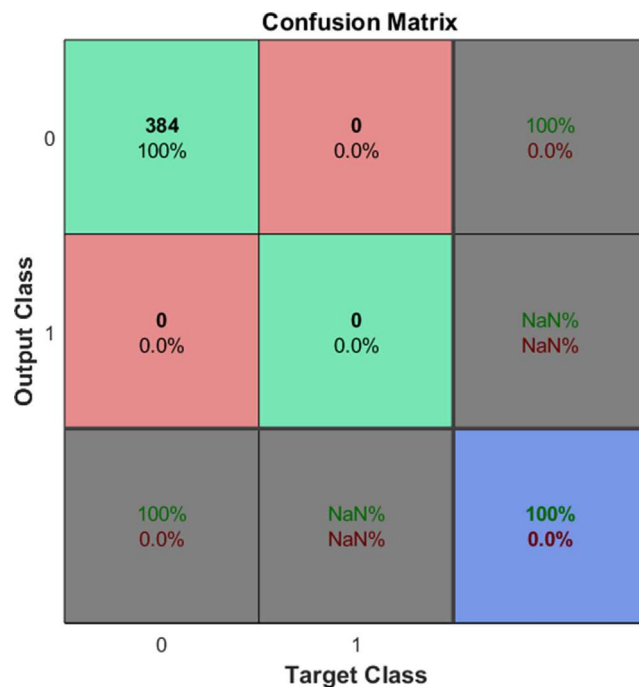


Fig. 20. A confusion matrix detailing the 'accuracy' of the decision algorithm for the second decision in the three time-slice problem for $U(F_r = 1) = -285$.

In this risk-based decision-framework, the costs are explicitly modelled and may be used to inform a preferential selection of classifier with regard to type I and type II errors.

Figs. 19 and 20 provide a breakdown of the 'decision accuracy' shown in Fig. 18 for the first and second decisions, respectively. Fig. 19 shows that, for the first decision, the algorithm was able to select the correct action in 332 of the 384 cases, yielding an overall 'decision accuracy' of 86.5%. Additionally, it can be seen that in 40 cases the 'perform maintenance' action

was selected incorrectly, and in 12 cases the ‘do nothing’ action was selected incorrectly. Comparing Figs. 18 and 19 reveals that all type I and type II errors occur during the first decision. Logically, Fig. 20 shows that the algorithm has a ‘decision accuracy’ of 100% with respect to the second decision. Moreover, Fig. 20 shows that all optimal decisions during the second time slice are ‘do nothing’. This result can be explained by considering two possibilities. During the first time step, if the algorithm has decided that maintenance is warranted, then, under the assumed transition model $P(\mathbf{H}_{t+1}|\mathbf{H}_t, d = 1)$, the structure is guaranteed to be in its undamaged health state in the second time step in which case further maintenance is unwarranted given $P(\mathbf{H}_{t+2}|\mathbf{H}_{t+1}, d = 0)$ and $U(F_T = 1) = -285$. Alternatively, if, during the first time step, the algorithm deems ‘do nothing’ to be the optimal decision, then the degradation forecast over two time steps is not sufficient to warrant a ‘perform maintenance’ action in the second time step; however, this decision, of course, is dependent on the utilities specified within the model.

Fig. 21 shows the number of state transitions an initial undamaged state may go through before a ‘perform maintenance’ action is decided when the cost of the failure event and cost of maintenance are varied and all other utilities remain as specified in Tables 1 and 2. It can be seen that the number of transitions decays exponentially as the cost of failure increases. This is in accordance with the intuitive understanding that as the cost of failure tends to infinity, the number of transitions before maintenance should be decided asymptotically approaches zero. It can also be seen from Fig. 21 that for a given failure event cost the time until maintenance decreases with cost of maintenance. It should also be noted that, logically, if the cost of maintenance exceeds the cost of failure, the structure will be allowed to operate until failure.

To investigate the influence of the cost of failure upon the overall ‘accuracy’ of the decision algorithm, the decision process used to produce Fig. 18 was repeated for varying $U(F_T = 1)$. It should be noted that the utility of ‘perform maintenance’ action was fixed at $U(d = 1) = -100$. In addition to varying $U(F_T = 1)$, the decision algorithm was executed assuming a uniform distribution over the health states targeted by the classifier rather than the distribution as predicted by the classifier. Fig. 22 shows how the ‘decision accuracy’ of each algorithm varies with the cost of the failure event.

Fig. 22 shows that the decision algorithms using the classifier and the uniform distribution assumption all have perfect ‘accuracy’ when $U(F_T = 1) = -100$. This result can be attributed to the fact that when the cost of maintenance is less than or equal to the cost of failure, the optimal decision is always ‘perform maintenance’, therefore, the algorithms all select this action independent of the information available regarding the health state of the structure.

It can be seen from Fig. 22 that, with the exception of the somewhat trivial case of $U(F_T = 1) = -100$, the decision algorithm utilising the classifier to incorporate probabilistic information about the health state of the structure into the decision consistently performs better than assuming a uniform distribution over health states. This is to be expected as the use of the classifier allows for the identification of the undamaged condition and differing damage states. On the other hand, as assuming a uniform distribution is entirely ignorant of the health state the selected decision will be invariant for a given failure event cost. Therefore, even when the structure is undamaged, assuming a uniform distribution will result in a maintenance action to be selected whilst the failure event cost is sufficiently high.

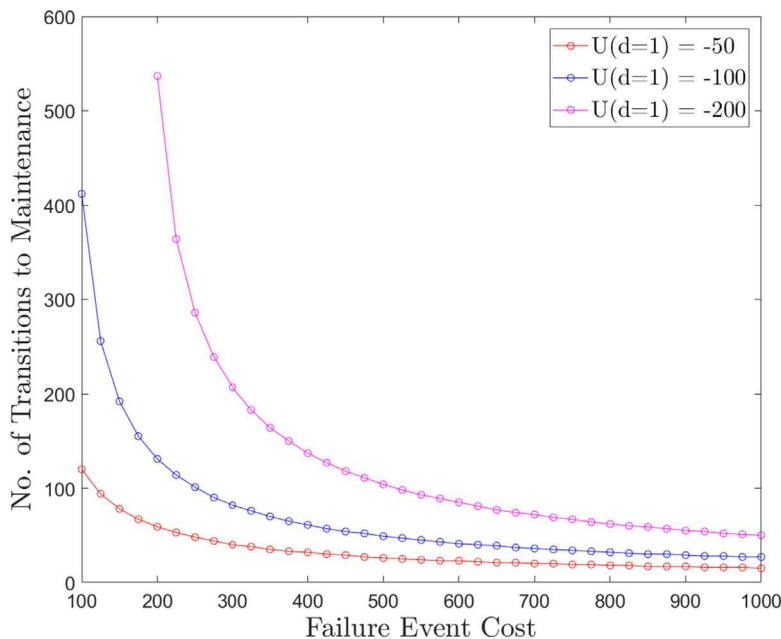


Fig. 21. The variation in the number of state transitions the undamaged structure will go through before maintenance is decided as a function of the failure event cost and maintenance cost. Failure event cost is defined as $-U(F_T = 1)$.

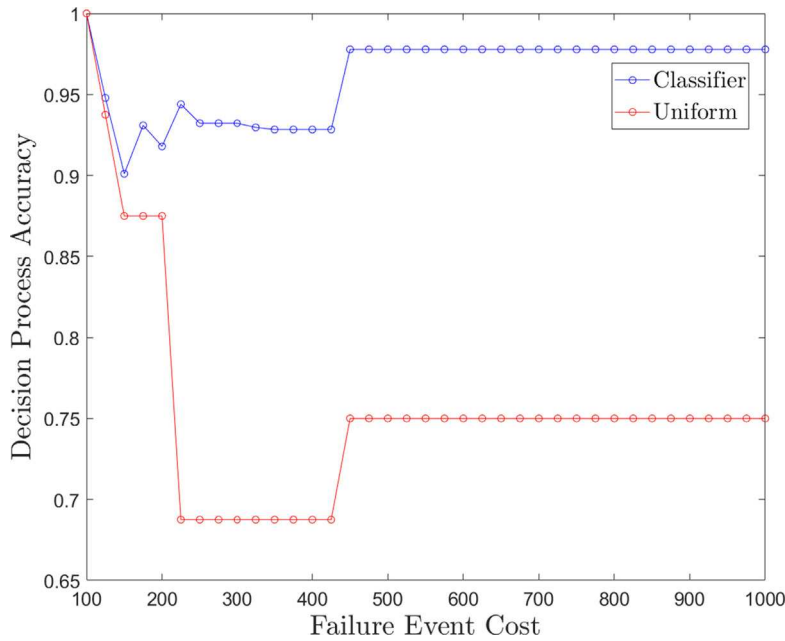


Fig. 22. A comparison of the 'accuracy' of the decision process as a function of the cost of the failure event when the health state is inferred using the statistical classifier and a uniform distribution over health states is assumed. Failure event cost is defined as $-U(F_T = 1)$.

The 'accuracies' of all three algorithms follow roughly the same trend for the range of failure costs shown in Fig. 22. For $U(F_T = 1) \leq -450$ the accuracy is constant. This is likely because, due to the high cost of failure, given the perfect knowledge of the health state, the optimal decision for all damage cases other than undamaged are 'perform maintenance'. For the algorithm utilising the classifier, misclassification of damage location does not influence the decided action in this cost range. As previously mentioned, the uniform assumption is ignorant of the health state and the decided actions are invariant in the range $U(F_T = 1) \leq -450$.

In Fig. 22, a lower 'decision accuracy' is seen in the failure cost range $-450 < U(F_T = 1) \leq -225$. This is expected to be because some of the optimal decisions given perfect information of the health for certain damage locations to be 'do nothing'. The uncertainty in health state reflected in the classifier and uniform distribution algorithms cause the 'perform maintenance' action to be deemed appropriate.

In the failure cost range $U(F_T = 1) \geq -200$ an increase in accuracy is seen for the algorithm assuming a uniform distribution over the health states. This is a result of the health state invariant decision being 'do nothing'; in this failure cost range this assumption is able to correctly decide actions for the undamaged cases and the less severe damage locations. The 'decision accuracy' of the algorithm employing the classifier fluctuates in this range, this may be due to the decided actions being sensitive to the distribution of uncertainties over the health states.

8. Discussion and conclusion

8.1. Summary

The framework described and demonstrated in the current paper provides an approach to risk-based decision-making in the context of SHM. Decision-making is facilitated through the inclusion of aspects of PRA such as fault tree modelling and risk, thereby allowing for the comparison of actions and the identification of a strategy that maximises expected utility.

The PRA paradigm currently practiced in industries such as aerospace and nuclear provides a basis for the formalisation of the operational evaluation procedure. Organising the information specifying the structure and monitoring system in a database will assist with ensuring all the necessary information required for subsequent stages is acquired and it will also provide a structured method for the retrieval of applicable information at each stage.

The fault tree development process of PRA provides the key novelty of this risk-based approach to SHM. Firstly, it facilitates the definition of key failure modes of interest and provides a structured method for identifying pertinent components whose health states should be targeted by a statistical classifier. The size of decision space for any given structure in the context of SHM is vast and an intimidating problem to begin addressing. By targeting selected failure modes of interest for a structure and modelling them as fault trees, the scope of the decision-maker may be limited, thereby making the problem more approachable; additional failure modes may subsequently be incorporated as an SHM system is further developed/expanded. Mapping the fault trees into Bayesian networks enables the framework to retain information regarding

the uncertainties in the health states thereby allowing robustness in the decision-making. Moreover, intermediate nodes within the Bayesian network representation of a fault tree may be queried, yielding marginal distributions that provide information about the probability of damage within components and/or substructures. This information may be utilised to guide inspection and maintenance engineers to specific locations, potentially saving time and reducing the cost of the actions, particularly for larger structures.

Although a relatively simple case study – specifically a truss – is used in the current paper for illustrative purposes, the probabilistic risk-based framework may be extended straightforwardly to many of the other structural types that exist in mechanical or civil infrastructure. This is conditional on there being sufficient information available in the SHM operational evaluation stage that health states of interest and key failure modes may be identified and precisely defined with respect to the hierarchical representation of the structure. Furthermore, the framework is not specific to the sensor and data type used in the case study; in fact, for certain SHM applications, where information of the dynamics is required, different sensors and data types will be a necessity. The framework may be applied in such cases through the actions taken during the data acquisition, feature extraction and statistical modelling stages of the risk-based SHM paradigm.

8.2. Open challenges

Whilst the framework presented addresses some of the problems surrounding the SHM decision process, there remain a number of challenges. One challenge, that has been widely acknowledged in the SHM community, is that data from the damage states of interest for a structure are seldom available prior to the implementation of an SHM system. This poses an issue in the development of the classifiers on which the decision process is highly dependent and a choice must be made regarding the approach to the statistical modelling. One option is to take a model-driven approach [24] that utilises outputs from physics-based models of the structure in its damage states of interest to learn a classifier in a supervised manner pre-implementation of the SHM system. Subsequently, the classifier can be continuously updated and validated with data obtained during the monitoring campaign. Alternatively, a semi-supervised approach can be taken in which a clustering algorithm is applied to the data acquired throughout the monitoring campaign. Clusters are attributed damage state labels through the incorporation of labelled data into the clustering algorithm [25]; damage state labels for data points may be obtained through inspection of the structure [26].

The results presented previously show that the performance of a probabilistic risk-based decision algorithm is dependent on the available information regarding the health state of the structure. As demonstrated, gains with regards to utility can be made in the absence of high-accuracy classifiers provided uncertainty is accounted for. This provides motivation for moving towards the use of true probabilistic classifiers in the context of SHM, be they discriminative (such as relevance vector machines (RVMs) [27]) or generative (such as GMMs [28]).

In addition to being dependent on the statistical classifier used, the optimality of decisions is highly contingent on the appropriateness of the transition model used; if the degradation of the structure is not accurately modelled, erroneous actions may be taken. Facing a similar issue to the statistical modelling process, oftentimes, data describing the transitions between the health states of interest are not held *a priori*. Again, one is faced with the choice of taking a model-driven approach involving the simulation of the degradation, or a data-driven approach that utilises data obtained during the monitoring campaign. The development of data-driven transition models, or validation of model-driven transition models is an awkward problem. Due to the fact that one is performing interventions on the structure during operation, information on transitions between health states is regularly censored, meaning the quantities of data spanning all state transitions of interest required for developing/validating transitions models may never be acquired, and particularly troublesome for one-off/bespoke structures. This problem is left as an open research question.

To ensure the desired performance of the decision algorithm, a vital stage in the risk-based framework is to assign utilities/costs to failure events and actions. Currently, within the literature there is no formal approach to how these values should be elicited, nor is there a consensus on how the risk preferences of an SHM decision-maker should be specified; should an agent be risk averse, risk neutral, or risk seeking? The issue at hand is one of both a technical and ethical nature, and whilst it will not be discussed in further detail in the current paper, it is highlighted to stimulate the conversations required for progress in the area of risk-informed decision-making for SHM.

8.3. Conclusion

In summary, a probabilistic risk-based framework for structural health monitoring was presented. Borrowing practices frequently used in probabilistic risk assessment, such as the use of fault trees to model system failures, the framework facilitates robust decision-making under uncertainty and provides advancements in the utility-optimal operation and maintenance of structures.

CRediT authorship contribution statement

A.J. Hughes: Conceptualization, Methodology, Software, Validation, Formal analysis, Writing - original draft, Writing - review & editing, Visualization. **R.J. Barthorpe:** Conceptualization, Supervision, Writing - review & editing. **N. Dervilis:**

Software, Methodology. **C.R. Farrar:** Conceptualization. **K. Worden:** Conceptualization, Funding acquisition, Investigation, Writing - review & editing.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors would like to acknowledge the support of the UK EPSRC via the Programme Grant EP/R006768/1. KW would also like to acknowledge support via the EPSRC Established Career Fellowship EP/R003625/1. The authors would like to thank Dr. Mark Bateman of EDF Energy for providing valuable discussions.

References

- [1] C.R. Farrar, K. Worden, *Structural Health Monitoring: A Machine Learning Perspective*, John Wiley & Sons Ltd, United Kingdom, 2013.
- [2] US Nuclear Regulatory Commission, PRA Procedure Guide, US Nuclear Regulatory Commission 1.
- [3] J.R. Fragola, G. Maggio, Space shuttle operational risk assessment, AIP Conference Proceedings 361 (May) (1996) 719–720, <https://doi.org/10.1063/1.49935>, URL: <http://aip.scitation.org/doi/abs/10.1063/1.49935>.
- [4] A. Hughes, K. Worden, R. Barthorpe, On an application of probabilistic risk assessment to structural health monitoring, in: Proceedings of the Twelfth International Workshop on Structural Health Monitoring, vol. 2, 2019, pp. 2400–2409, <https://doi.org/10.12783/shm2019/32376>.
- [5] E.B. Flynn, M.D. Todd, A Bayesian approach to optimal sensor placement for structural health monitoring with application to active sensing, *Mechanical Systems and Signal Processing* 24 (4) (2010) 891–903, <https://doi.org/10.1016/j.ymssp.2009.09.003>.
- [6] M. Vega, M. Todd, A variational bayesian neural network for structural health monitoring and cost-informed decision-making in miter gates, *Structural Health Monitoring* doi:10.1177/1475921720904543.
- [7] M. Gobatto, J.B. Kosmatka, J.P. Conte, A recursive Bayesian approach for fatigue damage prognosis: An experimental validation at the reliability component level, *Mechanical Systems and Signal Processing* 45 (2) (2014) 448–467.
- [8] J.S. Nielsen, Risk-based operation and maintenance of offshore wind turbines, Ph.D. thesis, Aalborg University, 2013, doi: 10.13052/rp-9788793102521.
- [9] M.K. Hovgaard, R. Brincker, Limited memory influence diagrams for structural damage detection decision-making, *Journal of Civil Structural Health Monitoring* 6 (2) (2016) 205–215, <https://doi.org/10.1007/s13349-016-0153-z>.
- [10] R. Schöbi, E.N. Chatzi, Maintenance planning using continuous-state partially observable Markov decision processes and non-linear action models processes and non-linear action models, *Structure and Infrastructure Engineering* 12 (8) (2016) 977–994, <https://doi.org/10.1080/15732479.2015.1076485>.
- [11] C. Li, S. Mahadevan, Y. Ling, S. Choe, L. Wang, Dynamic bayesian network for aircraft wing health monitoring digital twin, *AIAA Journal* 55 (2017) 1–12, <https://doi.org/10.2514/1.J055201>.
- [12] A. Rytter, Vibration based inspection of civil engineering structures, Phd thesis, Aalborg University, 1993.
- [13] T. Bedford, R. Cooke, *Probabilistic Risk Analysis: Foundations and Methods*, Cambridge University Press, Cambridge, United Kingdom, 2001.
- [14] US Nuclear Regulatory Commission, Probabilistic Risk Assessment (PRA), 2018, URL: <https://www.nrc.gov/about-nrc/regulatory/risk-informed/pr.html>.
- [15] L.E. Sucar, *Probabilistic Graphical Models: Principles and Applications*, Springer, London, 2015, doi: 10.1007/978-1-4471-6699-3.
- [16] J. Pearl, Fusion, propagation and structuring in belief networks, *Artificial Intelligence* 29 (3) (1986) 241–288.
- [17] U.B. Kjaerulff, A.L. Madsen, *Bayesian Networks and Influence Diagrams: A Guide to Construction and Analysis*, Springer, New York, 2008, doi: 10.1007/978-0-387-74101-7.
- [18] A. Bobbio, L. Portinale, M. Minichino, E. Ciancamerla, Improving the analysis of dependable systems by mapping Fault Trees into Bayesian Networks, *Reliability Engineering and System Safety* 71 (3) (2001) 249–260, [https://doi.org/10.1016/S0951-8320\(00\)00077-6](https://doi.org/10.1016/S0951-8320(00)00077-6).
- [19] S. Mahadevan, R. Zhang, N. Smith, Bayesian networks for system reliability reassessment, *Structural Safety* 23 (3) (2001) 231–251.
- [20] K. Worden, A.D. Ball, G.R. Tomlinson, Fault location in a framework structure using neural networks, *Smart Materials and Structures* 2 (3) (1993) 189–200.
- [21] M.F. Møller, A scaled conjugate gradient algorithm for fast supervised learning, *Neural Networks* 6 (4) (1993) 525–533, [https://doi.org/10.1016/S0893-6080\(05\)80056-5](https://doi.org/10.1016/S0893-6080(05)80056-5).
- [22] K.P. Murphy, *The Bayes Net Toolbox for Matlab*.
- [23] F. Jensen, F.V. Jensen, S.L. Dittmer, From influence diagrams to junction trees, *Uncertainty Proceedings* 10 (1994) 367–373, <https://doi.org/10.1016/B978-1-55860-332-5.50051-1>.
- [24] R.J. Barthorpe, On Model- and Data-based Approaches to Structural Health Monitoring A thesis submitted to the University of Sheffield for the degree of, Ph.D. thesis, The University of Sheffield, 2011.
- [25] L.A. Bull, T.J. Rogers, C. Wickramarachchi, E.J. Cross, K. Worden, N. Dervilis, Probabilistic active learning: an online framework for structural health monitoring, *Mechanical Systems and Signal Processing* 134.
- [26] T.J. Rogers, K. Worden, R. Fuentes, N. Dervilis, U.T. Tygesen, E.J. Cross, A Bayesian non-parametric clustering approach for semi-supervised Structural Health Monitoring, *Mechanical Systems and Signal Processing* 119 (2019) 100–119.
- [27] M.E. Tipping, Sparse Bayesian learning and the relevance vector machine, *Journal of Machine Learning Research* 1 (2001) 211–244.
- [28] C.M. Bishop, *Pattern Recognition and Machine Learning*, Springer Science + Business, LLC, New York, USA, 2006.