Acoustic Emission Source Location Using Bayesian Optimisation for a Composite Helicopter Blade

C. A. LINDLEY, M. R. JONES, T. A. DARDENO, R. S. MILLS, N. DERVILIS and K. WORDEN

ABSTRACT

It has been demonstrated that acoustic-emission (AE) based inspection of structures can offer advantages over other types of monitoring techniques in the detection of damage; namely, an increased sensitivity to damage, as well as an ability to localise its source. These advantages have resulted in the use of AE-based systems in Structural Health Monitoring (SHM) applications becoming an increasingly popular research topic. The strong correlation between features of AE-waves and the sources that generate them can be exploited to reveal information of the health-state of a structure in a non-invasive framework. This idea has been demonstrated successfully in the early detection of crack-growth, nucleation, corrosion, dislocation slips, and cavitation, among others. There are, however, numerous challenges associated with the analysis of AE-data. Some of these challenges include sensor placement, calibration, operational variations, sources of noise, and manipulation of large-size datasets. These challenges must be addressed carefully if one wishes to implement an AE-based system in SHM applications. Although the analysis of AE-data can be complex, its advantages can be meaningful in the development of an optimised maintenance schedule, reducing costs associated with unnecessary repairs, or by preventing potentially catastrophic outcomes from overlooking an emerging defect that can be detrimental to the safe operation of the structure. In this paper, AE data recorded from a full-scale helicopter blade are processed to develop a probabilistic model designed to predict the location of potential sources of damage. The following work addresses some of the challenges related to the construction of ΔT maps and proposes a novel strategy for identifying optimal sampling points, eliminating the need for extensive data collection for training.

INTRODUCTION

A research topic of increasing popularity in *Structural Health Monitoring* (SHM), is the use of *Acoustic Emission* (AE)-based techniques for damage detection. An AE is a sudden release of strain energy in an elastic medium [1], resulting in the propagation of very low amplitude waves characterised by frequencies in the ultrasonic range. One of

Christopher A. Lindley, PhD Student, Email: c.a.lindley@sheffield.ac.uk. Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, Mappin Street, Sheffield S1 3JD, UK

the main reasons these waves are of interest in SHM, is because damage in structures will promote the generation of AE activity [2]. Therefore, monitoring the progression of features extracted from AE-waves can be used to determine the health-state of the structure. Several promising results have been demonstrated in the literature when implementing AE-based monitoring techniques, with some examples including their application in bearings [3, 4], gears [5] and aerospace structures [6, 7].

This paper, however, is concerned with another attractive advantage of implementing AE-based techniques; that is, the possibility to localise the source of damage. Having multiple AE-sensors acting simultaneously, one can determine the location of an emerging defect by measuring the *Time of Arrival* (ToA) of a wave to each sensor.

A traditional approach to AE source localisation is based on triangulation, where the source is determined to be at the intersection of the contours resulting from calculating the differences in ToA (ΔT) from multiple sensor-pairs [8, 9]. Localisation may also be approached as an optimisation problem, where the origin of a source is estimated by minimising the difference between the recorded ΔT and values calculated at candidate locations.

Although these methods have been demonstrated to work successfully, there are a few associated challenges that can be arduous to overcome. For example, the propagating path of an emitted wave is an important factor to consider, as localisation becomes increasingly more difficult when dealing with composites, or when obstructions are present in the propagating path of the wave. Some progress has been achieved with anisotropic materials [10, 11], but it may still be challenging to analytically determine how the wave travels through the medium.

Data-based techniques have been developed for this type of application by instead learning statistical models that are uniquely concerned with the ΔT measurements [12]. The idea behind this approach is to artificially construct a map of ΔT values at various locations, to then match the ΔT from a real AE to those on the map. A probabilistic framework for this technique has also been explored and implemented successfully in [13, 14].

The problem is that ΔT maps can be laborious to produce. A vast collection of data may be required to cover the entire surface where damage might be expected. The ΔT maps must also be fine enough for accurate predictions, considerably increasing the size of the dataset needed for the learning process. In this paper, the efficient construction of a ΔT map is therefore explored, where a Bayesian optimisation approach is used for an optimal selection of sampling points. The subsequent sections detail the experimental procedure undertaken for data collection, outline the employed modelling strategies, and subsequently discuss the obtained results from the proposed methods.

EXPERIMENTAL PROCEDURE

The experimental set-up studied here was designed for a fatigue test of the helicopter blade. An eccentrically-loaded motor attached to the blade was tuned to have the blade resonate in its second mode of vibration, to promote the extension of fracture somewhere near the root.

The detection of the early onset of fracture in fatigue is possible with the use of AE-based monitoring systems [7, 15]. However, the effectiveness of monitoring the pro-

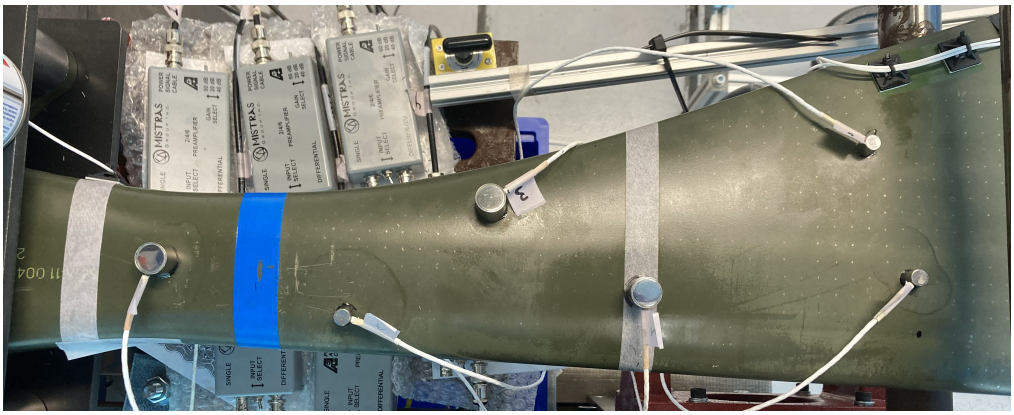

Figure 1. Helicopter blade and AE sensor configuration.

gression of fracture during the fatigue test relies on correctly characterising the detected AE waves. By having the monitoring system also locate the (possibly) various sources of AE, one can identify and isolate those deriving from damage. This consideration is important for correlating the features extracted from AE waves with the potential failure modes of the structure, thus highlighting the need for a robust localisation system. In [6, 16], for example, localisation techniques are exploited to directly characterise AE activity emerging near a developing crack.

The configuration of AE sensors used here is shown in Figure 1, where six sensors can be seen attached near the root of the blade. This set of sensors comprised three *WD differential* and three *Micro* 30*D differential Mistras* AE sensors. All sensors were pre-amplified with a *2/4/6 Switch Selectable Gain* set to $20dB$. The data acquisition system used was a *Mistras Micro-II compact PCI AE Chassis*.

In order to develop a localisation monitoring system, it was first necessary to collect enough data for training the model. The training set was therefore gathered by measuring the ΔT values of waves propagating through the blade. By knowing the location of the source beforehand, one can record the time it takes for the wave to trigger each sensor from its origin. The procedure was then followed with the construction of the ΔT maps for each sensor-pair combination. In this case, 15 unique sensor-pair combinations were possible.

A standard way to simulate realistic broad-band AE waves is by breaking pencil-leads on the surface of the blade. In this way, one can choose the location of the source using a *Hsu-Nielsen* device [17]. A collection of AEs can then be artificially generated at different locations and the respective ToA recorded to gather a training set. The challenge with this approach is on deciding where to sample on the blade. Intuitively, a grid could be placed over the blade with uniformly-spaced sample-points. The spacing between sampling points will determine the precision of the localisation algorithm, and thus a finer grid might be preferred in most applications. Unfortunately, there is the practical expense of having to sample from an inconveniently large number of points. The trade-off between accuracy and cost of the implementation will be application specific, warranting the need for an autonomous solution that can provide a balance between both.

Given that these samples had to be carried out by hand, a compromise was made in the fineness of the grid by having a sample-points spaced by $10mm$. Nevertheless,

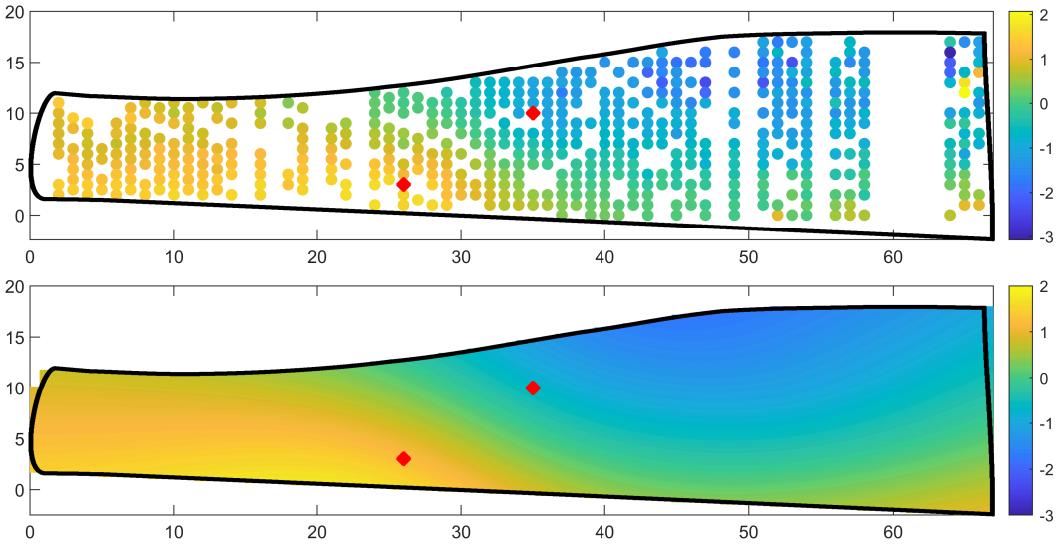

Figure 2. (Top) ΔT measurements recorded from sensor-pair 3-4. (Bottom) Inferred GP posterior mean. Red dots indicate the locations of sensors 3 (right) and 4 (left).

additional sample-points were included in between layers nearest to the root, since the fatigue test was designed to promote the extension of fracture in this area. The grid was designed on a graph-paper and was then wrapped over the blade to mark the locations of the sampling points on its surface. These points can be seen faintly in Figure 1.

Having established a sampling grid, the Hsu-Nielsen device was used to generate an AE-wave at each location. The waves were recorded and stored to then determine their true ToA at each sensor. This last step was necessary, given that the data acquisition system only places a time-stamp at the instance in which a threshold crossing occurs, which can happen even after the actual onset of the wave is picked up by the sensor(s). To account for this consideration, the true ToA of the recorded AE-waves can be determined based on the *Akaike Information Criterion* (AIC) [18]. The AIC-picker technique can be defined by the following equation,

$$AIC(t) = t\log(\mathbb{V}(R_w[1:t])) + (T-t-1)\log(\mathbb{V}(R_w[t:T])) \tag{1}$$

where at each time step t, the entropy is calculated before and after each time step. The entropy of the windowed signal, R_w, will return a minimum when the proportion of the signal up to the time step t is uniquely composed of uncorrelated noise, happening to also correspond to the onset of the AE-wave. The indexing in R_w indicate the range of values used to calculate the variance.

The results of the difference in time of arrival are shown in Figure 2 (top) for the sensor-pair 3-4. Some areas of the blade could not be sampled because of the sensors themselves occupying the space, or the limited accessibility because of the test-rig configuration, resulting in some sparsity that can be seen in the results. A somewhat gradual transition can be noted along the blade, which is to be expected in an ideal scenario, but a few inevitable exceptions can also be observed that appear in dispute with this trend. These "anomalous" data could be the result of noise in the measurements, mainly attributed to the complexity of the propagating paths in the blade.

A schematic of the internal structure of the blade was not available, but it is known to be made of a carbon composite with a complex internal honeycomb structure. The propagating paths could have therefore resulted in complex wave dispersion and/or attenuation that made it difficult to discern background noise from actual AE activity. In particular, waves originating furthest from a sensor were the most cumbersome to capture. After curation, the dataset comprised a total of 595 sampled points, and then was split into a training and test set, where 50 random sample points were allocated to the test set.

LOCALISATION STRATEGY

The localisation strategy used in this work follows the techniques explored by Jones et al. [19]. This approach is based on a regression model that maps a location on the blade to a corresponding difference in ToA, i.e. $(x, y) \rightarrow \Delta T$. Given the data gathered from the lead-breaks, it is possible to construct a model capable of returning the probability of observing a ΔT value at a new location (x_*, y_*). In this case, a *Gaussian Process* (GP) regression model [20] was chosen to predict ΔT values over the continuous space covered by the grid. Fundamentally, the GP attempts to model functions of the form,

$$t = f(x) + \epsilon, \qquad \epsilon \sim \mathcal{N}(0, \sigma_n^2) \tag{2}$$

where t is a target value for a given input x, $f(x)$ is some function modelled by the GP and ϵ corresponds additive Gaussian noise with zero mean and variance σ_n^2. The GP is defined by two sets of input points x and x' as,

$$f(x) \sim \mathcal{GP}(m(x), k(x, x')) \tag{3}$$

Making predictions for a new test point t_* involves first taking the joint Gaussian distribution over all points of interest,

$$\begin{pmatrix} \mathbf{t} \\ f_* \end{pmatrix} \sim \mathcal{N}\left[\begin{pmatrix} m(X) \\ m(x_*) \end{pmatrix}, \begin{pmatrix} K(X,X) + \sigma_n^2 & K(X, x_*) \\ K(x_*, X) & K(x_*, x_*) \end{pmatrix} \right] \tag{4}$$

where X is used to denote a matrix formed of all the N training inputs and $\mathbf{t}$ is a vector of the corresponding set of N target observations. The distribution over the set of latent functions $\mathbf{f}_*$, conditioned on the observations, can be expressed as,

$$p(\mathbf{f}_*|\mathbf{t}) = \mathcal{N}(\mathbb{E}[\mathbf{f}_*], \mathbb{V}[\mathbf{f}_*]) \tag{5}$$

where the posterior mean $\mathbf{E}[\mathbf{t}_*]$ and variance $\mathbf{V}[\mathbf{t}_*]$ are defined as,

$$\mathbb{E}[\mathbf{f}_*] = m(X_*) + K(X_*, X)(K(X,X) + \sigma_n^2 \mathbb{I})^{-1}\mathbf{t} \tag{6}$$

$$\mathbb{V}[\mathbf{f}_*] = K(X_*, X_*) - K(X_*, X)(K(X,X) + \sigma_n^2 \mathbb{I})^{-1} K(X, X_*) \tag{7}$$

As before, the notation X_* now includes all the test inputs. The covariance (kernel) function must be chosen to fully specify the GP. This function defines the similarity between any two sets of input points, resulting in a symmetric semi-positive square covariance matrix K. An example of a commonly-used covariance function would be the *Matérn* $3/2$ function,

$$k(x, x') = \sigma_f^2 \left(1 + \frac{\sqrt{3}r}{l}\right) \exp\left(-\frac{\sqrt{3}r}{l}\right), \qquad r = ||x - x'|| \tag{8}$$

where σ_f^2 and l are the *hyperparameters* of the GP and characterise the signal variance and the length scale of the functional form of $f(x)$, respectively. This particular covariance function is recommended when modelling physical processes [21]. Here, the *Quantum-Behaved Particle Swarm* (QBPS) optimisation algorithm [22] was employed to find the optimised hyperparameters that maximise the marginal likelihood.

Implementing the GP with the results presented in the previous section, returns an interpolated mapping over the space covered by the grid, as shown in Figure 2 (bottom). By repeating the implementation of the GP on all 15 sets of results (corresponding to each sensor-pair), one can make use of these models to pinpoint the most likely location from which an AE-wave may originate.

However, in this form, the predictive distribution cannot directly provide the location of a new AE wave. The location can be determined by instead calculating the likelihood of observing the set of ΔT values for each sensor-pair, given the functional mapping $(x, y) \rightarrow \Delta T$ learned from the training set. For the set of $J = 15$ models, the log-likelihood for each model m_j can be assessed as,

$$\log p(t_{*,j}|D, (x_*, y_*), m_j) = -\frac{1}{2}\log \mathbb{V}[\mathbf{f}_{*,j}] - \frac{(t_{*,j} - \mathbb{E}[\mathbf{f}_{*,j}])^2}{2\mathbb{V}[\mathbf{f}_{*,j}]} - \frac{1}{2}\log 2\pi \tag{9}$$

where $D = (\mathbf{x}, \mathbf{y})$ denotes the location of all data-points in the training set, and $t_{*,j}$ is the predicted ΔT at the new location (x_*, y_*), given by the j^{th} sensor-pair. The log-likelihood maps produced by assessing equation (9) will display contours of the most probable locations of $t_{*,j}$. A unique solution is generally found at the intersection of the contours evaluated for each sensor-pair model. Mathematically, this outcome can be derived by marginalising each model. That is,

$$p(\boldsymbol{t}_*|D, (x_*, y_*)) = \sum_{j=1}^{J} p(t_{*,j}|D, (x_*, y_*), m_j)p(m_j) \tag{10}$$

If each model is assumed to have equal importance, the marginalised probability distribution is given by,

$$p(\boldsymbol{t}_*|D, (x_*, y_*)) = \sum_{j=1}^{J} p(t_{*,j}|D, (x_*, y_*), m_j) \tag{11}$$

which is simply the sum of the likelihood maps for each sensor-pair. The most likely location will therefore be given by the point (x_*, y_*) that maximises equation (11). The outcome of this approach for an arbitrary test point is shown in Figure 3.

BAYESIAN OPTIMISATION STRATEGY

As mentioned above, the sampling grid is subject to some practical limitations. In this section, the idea of developing an optimum grid is explored employing *Bayesian*

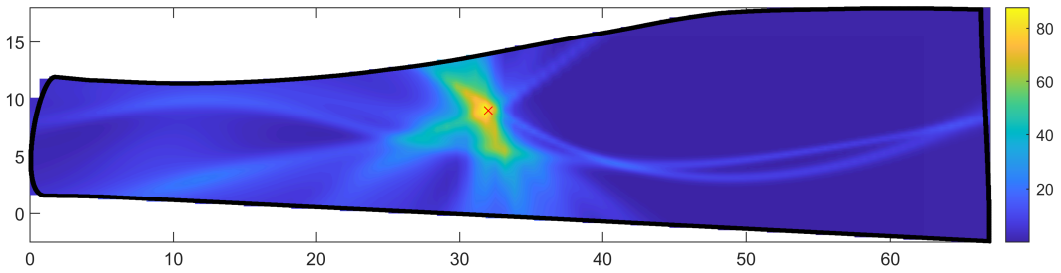

Figure 3. Prediction of the GP model for an arbitrary test point. The red x marks the true location of the AE source.

Optimisation (BO) [23]. In a BO scheme, an objective function is optimised by finding its local maximum (or minimum) with as few parameter proposals as possible. This method does not rely on gradients and can be an attractive alternative in applications where the objective function cannot be easily determined or is expensive to compute.

The objective function is approximated with a surrogate function, which is usually chosen to be a GP. A new proposal reveals its corresponding output of the objective function, followed by conditioning the surrogate GP with the revealed observations. Deciding on the best proposal is determined by an *acquisition function*. A variety of acquisition functions exist, but the commonly used *Expected Improvement* (EI) [23] is adopted here. The EI is defined as,

$$EI(x) = \begin{cases} (\mu(x) - f(x^+) - \eta)\Phi(Z) + \sigma(x)\phi(Z) & \text{if } \sigma(x) > 0 \\ 0 & \text{if } \sigma(x) = 0 \end{cases} \tag{12}$$

where $\mu(x)$ and $\sigma(x)$ are the mean and standard deviation of the surrogate GP posterior, respectively. $\Phi(Z)$ and $\phi(Z)$ correspond to the *Cumulative Density Fucntion* (CDF) and *Probability Density Function* (PDF) of the standard normal distribution, evaluated with respect to $Z = (\mu(x) - f(x^+) - \eta)/\sigma(x)$, respectively.

Acquisition functions suggest new proposals by balancing the exploitation and exploration of the search-space. In short, exploitation refers to sampling more from areas that improve the output of the objective function, while exploration promotes sampling from areas of high uncertainty. The parameter η in equation (12) can be adjusted to control the amount of exploration.

Although used for finding optimal parameters, the BO can be adopted to select the best points to sample from in the ΔT grid. The idea here is to minimise the number of samples one would need to construct a grid simple enough to determine the location of new AE-waves. Simultaneously, this approach would also avoid having to decide on the density of the grid, since new samples are proposed sequentially in an active setting. This idea has been explored in [24], where a BO scheme is used to promote sampling from damaged areas in specimens for an autonomous ultrasonic inspection.

What remains is to decide on the objective function for this scenario. An ideal objective function may be one that promotes sampling in areas that yield the highest gain in accuracy while allowing for some exploration in areas of high uncertainty in the blade. The objective function used here is one constructed by calculating the reduction in the *Root-Mean-Square Error* (RMSE) that a single sample-point in the training set provides

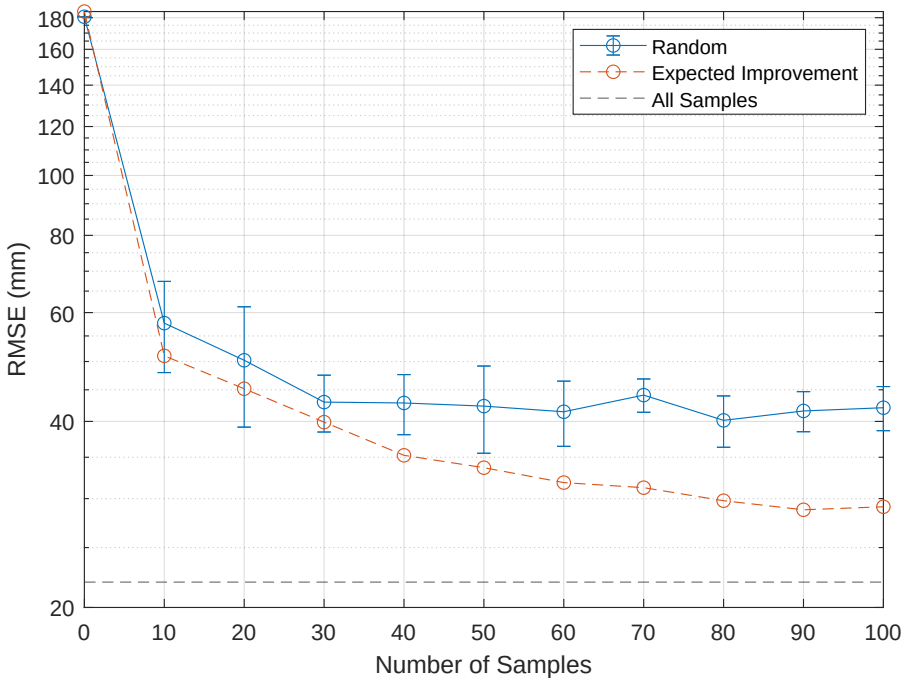

Figure 4. RMSE calculated with data sampled from a random process (blue) and by using the EI acquisition function (red). The RMSE calculated from a complete dataset has been included as a reference (black dashed line).

to the localisation of the AE-waves in the test set. The RMSE is expressed as,

$$RMSE = \sqrt{\frac{\sum_{x,y}((x,y)_{pred} - (x,y)_{true})^2}{N_{test}}} \tag{13}$$

where N_{test} is the number of samples in the test set. If assessed for each point on the grid at a time, the surrogate objective would look like that shown in Figure 5 (top). One should note that this objective corresponds to the final form that the surrogate would achieve if all sample points were taken into consideration. In practice, this objective function is approximated, as samples are sequentially introduced.

RESULTS AND DISCUSSION

To demonstrate the implementation of the BO strategy for an efficient ΔT map construction, a scenario is considered in which limiting resources only allow for 100 samples to be taken. These samples would follow on the test set already collected at 50 random locations within the borders of the grid. The results of the sequential sampling are shown in Figure 4. A series of random-sampling strategies were also performed to demonstrate the benefits of choosing an optimal (under the BO) collection of sampling points. The RMSE calculated from having sampled all training data-points has also been included as a reference.

A few observations are worth addressing from these results. Firstly, the BO strategy appeared to have found a sequence of samples that returns a reduction in error at a greater

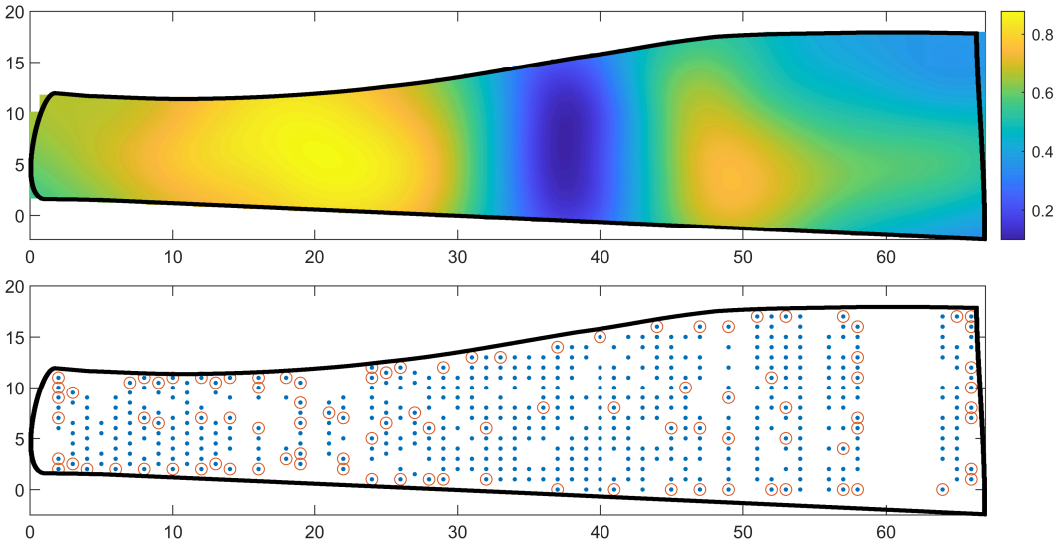

Figure 5. (Top) Objective function over the search-space. (Bottom) Red circles indicating the sampled locations proposed by the EI acquisition function.

rate than the random sequence. It may be possible to find a better combination of points at random but at the high risk of having to reiterate the test if otherwise. The results from the random test also appear to show that some data-points can negatively affect the performance of the model. This outcome could be explained by the "anomalous" measurements described above. It may be necessary to repeat this study but with a "clean" dataset gathered from a simpler structure, such as an aluminium plate.

Another highlight can be discussed when observing the candidate points selected for sampling in Figure 5 (bottom). A higher concentration of samples can be seen to have been taken around the global maximum of the objective function. It is important to note, however, that the aim here deviates from the purpose of a conventional Bayesian optimiser in finding a unique solution that maximises the objective function. Although this was the intention of the proposed strategy, some substantial exploration had to be enforced by increasing the parameter η. This last step is important since the accuracy of the localisation model improves with a comprehensive representation of the search-space, which is achieved by gathering data covering the entirety of the blade. A balance of these aspects is nicely encapsulated by the acquisition function.

Finally, it is worth acknowledging that the proposed objective function is by no means the only option one could use in practice. In [24], for example, the objective function is based on novelty detection techniques used in SHM. The success of this approach may rely on the correct choice of the objective function, acquisition function, and correct tuning of the parameters involved. Deciding on these parameters will naturally depend on the structure and application at hand.

CONCLUDING REMARKS

As demonstrated in this paper, BO can be used for optimal sequential sampling of ΔT maps. The construction of these maps is a requisite in the implementation of prob-

abilistic techniques for damage localisation in SHM. The proposed strategy is based on an active scheme that proposes an optimal candidate point for sampling based on EI; therefore, reducing the number of necessary measurements needed for the construction of a robust localisation system.

Some considerations were mentioned; mainly, deciding on the best objective function and careful tuning of the corresponding parameters to balance exploitation and exploration of the studied surface.

Future work will involve using the developed localisation techniques to identify fracture extension during the fatigue test. Additionally, the presented ideas will be explored further with the implementation of different objective functions, or a combination of them in a multi-objective optimisation scheme.

ACKNOWLEDGMENT

The authors gratefully acknowledge the support of the UK Engineering and Physical Sciences Research Council (EPSRC), via grant reference EP/W005816/1, as well as grant references EP/R003645/1 and EP/R004900/1. For the purpose of open access, the author(s) has/have applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising. This research made use of The Laboratory for Verification and Validation (LVV), which was funded by the EPSRC (via EP/J013714/1 and EP/N010884/1), the European Regional Development Fund (ERDF), and the University of Sheffield. The authors would like to extend special thanks to Michael Dutchman, for helping set up the experiments.

REFERENCES

1. Kaiser, J. 1950. *An investigation on the occurrence of noises in tensile tests or a study of acoustic phenomena in tensile tests*, Ph.D. thesis, Tech Hochsch.
2. Ono, K. 2011. "Application of acoustic emission for structure diagnosis," *Diagnostyka*:3–18.
3. Choudhury, A. and N. Tandon. 2000. "Application of acoustic emission technique for the detection of defects in rolling element bearings," *Tribology International*, 33(1):39–45.
4. Poddar, S. and N. Tandon. 2019. "Detection of particle contamination in journal bearing using acoustic emission and vibration monitoring techniques," *Tribology International*, 134:154–164.
5. Feng, P., P. Borghesani, H. Chang, W. Smith, R. Randall, and Z. Peng. 2019. "Monitoring gear surface degradation using cyclostationarity of acoustic emission," *Mechanical Systems and Signal Processing*, 131:199–221.
6. Holford, K., R. Pullin, S. Evans, M. Eaton, J. Hensman, and K. Worden. 2009. "Acoustic emission for monitoring aircraft structures," *Proceedings of the Institution of Mechanical Engineers Part G Journal of Aerospace Engineering*, 223.
7. Hensman, J., K. Worden, M. Eaton, R. Pullin, K. Holford, and S. Evans. 2011. "Spatial scanning for anomaly detection in acoustic emission testing of an aerospace structure," *Mechanical Systems and Signal Processing*, 25(7):2462–2474.
8. Kundu, T., H. Nakatani, and N. Takeda. 2012. "Acoustic source localization in anisotropic plates," *Ultrasonics*, 52(6):740–746.

9. Tobias, A. 1976. “Acoustic-emission source location in two dimensions by an array of three sensors,” *Non-Destructive Testing*, 9(1):9–12.
10. Yamada, H., Y. Mizutani, H. Nishino, M. Takemoto, and K. Ono. 2000. “Lamb wave source location of impact on anisotropic plates,” *Journal of Acoustic Emission*, 18.
11. Kundu, T., S. Das, S. Martin, and K. Jata. 2008. “Locating point of impact in anisotropic fiber reinforced composite plates,” *Ultrasonics*, 48(3):193–201.
12. Baxter, M., R. Pullin, K. Holford, and S. Evans. 2007. “Delta T source location for acoustic emission,” *Mechanical Systems and Signal Processing*, 21(3):1512–1520.
13. Hensman, J., R. Mills, S. Pierce, K. Worden, and M. Eaton. 2010. “Locating acoustic emission sources in complex structures using Gaussian processes,” *Mechanical Systems and Signal Processing*, 24(1):211–223.
14. Jones, M., T. Rogers, K. Worden, and E. Cross. 2022. “A Bayesian methodology for localising acoustic emission sources in complex structures,” *Mechanical Systems and Signal Processing*, 163:108143.
15. Lindley, C., M. Jones, T. Rogers, E. Cross, R. Dwyer-Joyce, N. Dervilis, and K. Worden. 2023, “A probabilistic approach for acoustic emission based monitoring techniques: with application to structural health monitoring,” .
16. Hensman, J., R. Pullin, M. Eaton, K. Worden, K. Holford, and S. Evans. 2009. “Detecting and identifying artificial acoustic emission signals in an industrial fatigue environment,” *Measurement Science and Technology*, 20(4):045101.
17. Hsu, N. 1977, “Acoustic emission simulator,” US Patent 4,018,084.
18. Kurz, J., C. Grosse, and H. Reinhardt. 2005. “Strategies for reliable automatic onset time picking of acoustic emissions and of ultrasound signals in concrete,” *Ultrasonics*, 43(7):538–546.
19. Jones, M. 2023. *On Novel Machine Learning Approaches for Acoustic Emission Source Localisation: A Probabilistic Perspective*, Ph.D. thesis, University of Sheffield.
20. Rasmussen, C. and C. Williams. 2005. *Gaussian Processes for Machine Learning*, The MIT Press.
21. Stein, M. 1999. *Interpolation of Spatial Data: Some Theory for Kriging*, Springer Series in Statistics, Springer New York.
22. Sun, J., B. Feng, and . X. W. 2004. “Particle swarm optimization with particles having quantum behavior,” in *Proceedings of the 2004 Congress on Evolutionary Computation*, vol. 1, pp. 325–331 Vol.1.
23. Donald, J., M. Schonlau, and W. Welch. 1998. “Efficient Global Optimization of Expensive Black-Box Functions,” *Journal of Global Optimization*, 13(4):455–492.
24. Fuentes, R., P. Gardner, C. Mineo, T. Rogers, S. Pierce, K. Worden, N. Dervilis, and E. Cross. 2020. “Autonomous ultrasonic inspection using Bayesian optimisation and robust outlier analysis,” *Mechanical Systems and Signal Processing*, 145:106897.