Non-destructive structural characterization of graphite components using mechanical resonance and deep learning

Paul Geimer, Andrew Delorey, and Rajarshi Bose

LA-UR-25-29538
September 26, 2025

Los Alamos
NATIONAL LABORATORY

Prepared for: U.S. Department of Energy, Office of Nuclear Energy

Prepared by: Paul Geimer, Scientist, MST-8
Andrew Delorey, Scientist, EES-16
Rajarshi Bose, Undergraduate Researcher, MST-8

Prepared on behalf of: DOE-Nuclear Energy Microreactor Program

Contents

Figures

Acknowledgements

This work is supported by the Microreactor Program of the US Department of Energy's Office of Nuclear Energy. Los Alamos National Laboratory is operated by Triad National Security, LLC, for the National Nuclear Security Administration of U.S. Department of Energy (Contract No. 89233218CNA000001). The authors would like to thank Aneesh Pawar and Anna Patelli at LANL for their contributions to data acquisition systems and modeling, respectively; as well as Erik Luther (LANL) and Peter Hosemann and Alexandros Spyromilios (UC Berkeley) for providing samples.

Abbreviations

DOE	Department of Energy
FEA	finite-element analysis
FNO	Fourier neural operator
GNN	graph neural network
LANL	Los Alamos National Laboratory
LDV	laser Doppler vibrometer
MRP	Microreactor Program
NEAMS	Nuclear Energy Advanced Modeling and Simulation
ORNL	Oak Ridge National Laboratory
SHM	structural health monitoring

SUMMARY

As compared to conventional nuclear reactors, microreactors have the potential to significantly reduce construction timelines and capital costs, decreasing the barriers for advanced nuclear reactor technologies. However, the lower power output of these microreactors (typically < 20 MWe) creates challenging economics if operation and maintenance costs cannot be sufficiently reduced. The compact size of these designs presents an opportunity for comprehensive in-situ structural health monitoring to provide real-time feedback in order to reduce operational costs associated with maintenance and downtime. Many microreactor concepts use graphite for both in-core neutron moderation and as a structural material, which has typically required some form of periodic and laborious inspection.

This report provides a description and assessment of recent work with graphite to couple acoustic-based experimental measurements and characterization with machine learning models to mature structural health monitoring capabilities and generate benefits for the nuclear microreactor industry. With resilient embedded sensors in development in other programs funded by the US Department of Energy's Office of Nuclear Energy and elsewhere, the work described herein builds upon previously funded efforts to mature non-destructive testing technology that relates measured vibrational signatures to structural changes, using a combination of new experimental measurements and machine learning processing.

Building on past successful demonstrations of predictive workflows to identify structural changes in a hexagonal stainless steel test article with excellent acoustic propagation, we first performed baseline characterization on graphite samples with canonical geometries to ensure compatibility and confidence in the applied techniques for a material with distinctly different mechanical properties. In contrast to efforts in previous years, we worked exclusively with unidirectional vibration data that is more comparable to those expected from the existing embedded sensor technologies which are suitable for deployment in a reactor setting. Established acoustic and modern machine-learning-based characterization approaches were applied to the resulting datasets from these simple geometries. Both approaches were found to be highly capable of detecting even small geometric irregularities amongst nominally identical samples. As such, we then moved to testing these approaches for detection of artificial local stress perturbations introduced into a more complex geometry: a hexagonal block with drilled holes.

A main outcome of this work is that a generalizable ML workflow can be used to detect and predict the characteristics of small artificial anomalies in a graphite component with a relevant geometry. While this work was performed using surficial vibration data, we expect the approach to be flexible and viable for other monitoring scenarios, such as those with different arrangements or types of sensor arrays. As compared to previously funded efforts, an existing ML workflow based on neural networks was enhanced through the addition of recently developed Fourier neural operators. As applied to previously collected and new vibration datasets, prediction accuracies of anomaly characterizations were greatly improved with minimal added computational cost. As trained on small durations of vibration data (tens of seconds) collected over a realistic number of locations, the model was able to reliably determine the presence of a subtle stress anomaly and begin to provide location estimates. Such an approach is likely to be viable for more relevant reactor damage scenarios for graphite components, such as progressive crack growth or creep.

1 Introduction

Microreactors are compact nuclear energy systems designed for versatility and portability, where extreme operating conditions are likely to subject their components to high temperatures, radiation, and fatigue. Ensuring safe and reliable operation in such conditions with minimal downtime for inspection would reduce operating costs to make the technology more economically viable. This provides motivation for the development and implementation of non-invasive techniques for early detection of structural defects, especially those which are scalable to provide a level of comprehensive monitoring that is not feasible within larger and more complex reactor designs.

Since the first artificial nuclear reactor (Chicago Pile-1), graphite has been a key reactor material due to its neutron moderating properties, low cost, and high thermal stability. Many commercial designs since have utilized graphite as a dual-purpose material which also provides structure within a core assembly. As a quasi-brittle material, reactor-grade graphite has been extensively studied to understand how damage can accumulate due to dimensional changes from the effects of anisotropic stress and radiation damage [1–5]. These studies have been used to develop 3D damage models, such as those predicting fracture propagation due to asymmetric neutron flux [6]. Graphite structural material is a common element of many proposed microreactor designs. However, inspection of graphite reactor components has typically required their physical removal and some form of visual or enhanced imaging. As this is not compatible with the remote and/or semi-automated operational use cases that have been proposed for microreactors, non-invasive technologies that can provide enhanced or early indications of structural changes within graphite components are desired.

In FY25, Instrumentation and Sensors work within the Technology Maturation technical area of the Microreactor Program (MRP) at Los Alamos National Laboratory (LANL) focused on non-destructive vibration-based inspection of bulk graphite samples for defect characterization, aiming to perform structural health monitoring (SHM) at the component level. We conducted experiments on graphite using a combination of experimental modal vibration testing and machine-learning (ML) data analyses to test and validate experimental conditions, progressing the approach towards potential deployment in future microreactor testbeds and designs. This work builds on techniques developed for comparable stainless-steel components, where previous MRP-funded work on damage characterization for an additively manufactured 316L stainless steel test article yielded ML predictions of artificially modified structural states that were up to 90% accurate [7]. Results presented within describe significant prediction improvements produced by small changes to an existing ML workflow and suggest compatibility of this workflow for SHM in graphite components.

2 Sample descriptions

Graphite samples procured for testing were all composed of high-strength isotropic material, comparable with that used in numerous operational and proposed reactor designs [8–10], making for valid vibrational and elastic responses. While the grades of graphite differed slightly across the tested samples, all had similar properties: fine-grained (typical grain sizes of 10 – 13 μm) and high density (1.82 – 1.85 g/cm^3), with typical porosities around 10% and Young's moduli of about 11 GPa.

Graphite test samples consisted of a set of eight 6-cm diameter machined spheres (Tokai grade G347) similar to those acoustically inspected in a recent study [11], a set of four 25.4 x 305 mm (1" x 12") cylindrical rods (Bar-Lo grade B325), and last—to enable testing on a relevant and complex geometry—a monolithic hexagonal block (Mersen grade 2124) with 19 equidistant and equal-diameter holes machined through its length [12].

3 Methods

3.1 Experimental Approach

Three primary acoustic measurement procedures were used for ML-aided characterization: resonant ultrasound spectroscopy (RUS), mode shape imaging, and time-domain response. In all procedures, vibration measurements were made using a Polytec fixed single-point VibroFlex QTec laser Doppler vibrometer (LDV; short wavelength infrared) coupled with a two-axis scanning mirror galvo system (Thorlabs) to enable the laser to be both precisely aimed and rapidly rastered across a sample for imaging surface displacements. We used a bare lead-zirconate-titanate piezoelectric transducer (2mm x 12.7 mm disc) as the vibration excitation source, in dry contact with each of the samples, with a 20x gain Tegam amplifier connected in-line to the function generator to ensure sufficient signal-to-noise across a broad range of input frequencies. For testing, the hex block was laid onto one of its long faces with a flat piece of open-cell foam between the sample and benchtop to isolate it from surrounding vibration sources. The spheres and rods were supported in a tripod configuration to enable repeatable boundary conditions between samples, with the transducer forming one leg of the tripod and rounded bolt heads for the other two. All supports were attached to the same small optical breadboard with an open-cell foam underlay.

RUS measurements were used to establish baseline elastic behavior for the sample under highly controlled conditions, with the response of the samples measured using a swept-sine excitation from the established Resonant Inspection and Transducer Acoustics (RITA) software. Identified peaks in the frequency spectrum were then confirmed to be resonance modes of the samples using experimental modal analysis, whereby one face of the sample was rapidly scanned by the laser while the sample was excited at a series of resonances frequencies, producing an image of the mode shape that could be compared to modeled predictions using finite-element analysis (FEA, COMSOL Multiphysics v6.2).

The time-domain vibration response of the samples under a steady-state noise excitation provided the raw data inputs for training the ML model to make predictions of structural condition. These datasets were collected using a custom LabView interface that synchronized excitation and measurement timings, enabling repeatable user-defined waveforms to be sent into the sample via the transducer, with the LDV measuring at discrete points on the outer surface of the samples. As in previously funded work on this topic, we used a white-noise signal for excitation to better replicate the broadband vibration sources that might be expected in a passive SHM deployment, where no purpose-built excitation source is required to perform a structural analysis. A typical measurement at a single location during sustained noise excitation generated an amplitude time series with five million data points, collected at 250 kHz over 20 seconds. The chosen sampling rate was partially informed by a simple sensitivity study on hex block data which showed that, all else being equal, the revised ML model used in this work (see Section 4) trained more slowly when given datasets containing higher frequencies.

Experimental work was designed to produce large quantities of vibrational data on pristine and defective graphite samples in order to provide enough training examples for an ML model to make successful predictions of defect characteristics on unseen data. Towards this goal, we first performed testing and elastic characterization to understand batch variability and refine testing protocols. This consisted of acquisition and analysis of vibration data obtained at a single, consistent point across a set of samples, as well as the aforementioned mode shape imaging as needed.

Upon completion of this initial testing, an experiment was designed to generate structural anomalies within the hex block. To better allow for repeated testing using a single sample, the anomaly introduced within block was recoverable/elastic in nature, rather than permanently damaging. The chosen method to accomplish this was an expanding, removable plug designed to generate repeatable and localized compressive stress, such as those which could be generated by swelling or anisotropic thermal expansion

(Figure 1). The plug was constructed from a 1.5-cm long rubber cylinder inserted onto the end of a threaded rod. Metal washers were positioned on each side of the cylinder. An aluminum tube sized to fit over the threaded rod transferred the compression of an external nut into the stopper, causing it to expand in place within the sample. Though only two discrete locations (near block edge, block middle) are discussed in this report, this assembly is capable of being precisely inserted and affixed at any depth within the sample. This setup enabled localized stresses to be applied internally to the hex block, leaving the outer surface unblemished for external measurement by the LDV. While the stresses applied by the plug to the internal surfaces were not directly quantified, it was made repeatable using a torque meter on the wrench used to tighten the external compressive nut.

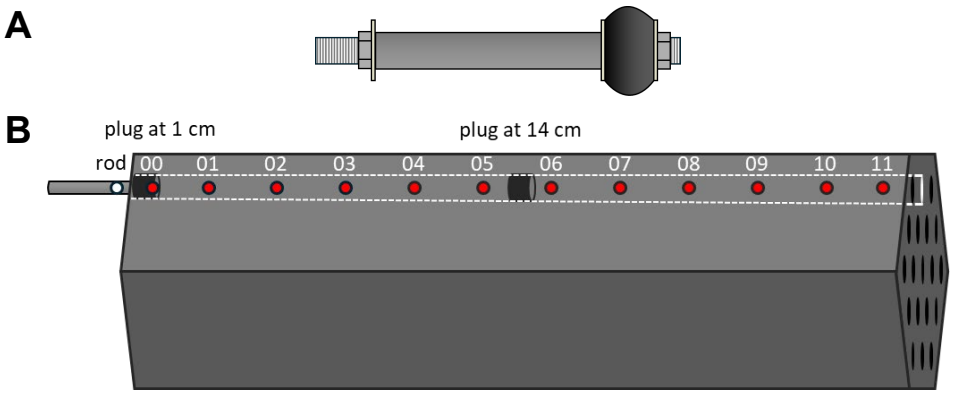

Figure 1. Diagram of experimental testing setup for hex block. (A) Insertable rubber plug assembly. (B) Plug and LDV measurement locations on the hex block. Diagrams not drawn to scale. Block geometry is comparable to those described in Trellue et al. [12].

An example of an as-recorded vibration time series and its spectral content is given in Figure 2, showing the steady-state response of the sample both in time and in frequency. Different locations were measured by horizontally translating the entire assembly in front of the LDV beam to ensure a consistent, normal angle between the beam and block surface. This prevented undesirable measurements of horizontal components of motion from a beam hitting the surface at a changing and oblique angle, which could unintentionally encode location information within the training data given to the ML model.

We note that most measured responses included several distinct resonances originating from the galvo system in the measurement setup rather than the samples themselves. Primarily located between 7 – 30 kHz, these resonances were distinguished from the sample resonances by (i) their consistent frequency and amplitude regardless of the laser location or excitation level of the transducer, and (ii) their broad peak widths that contrasted with the narrow peaks associated with the resonances of the samples. While these features would typically be viewed as undesirable spectral contamination for modal analysis, we chose to leave them in the dataset as an opportunity to perform a more realistic evaluation of our ML approach for real-world applications, where any number of nearby components or processes would be likely to introduce spectral features unrelated to the structural health of a component of interest.

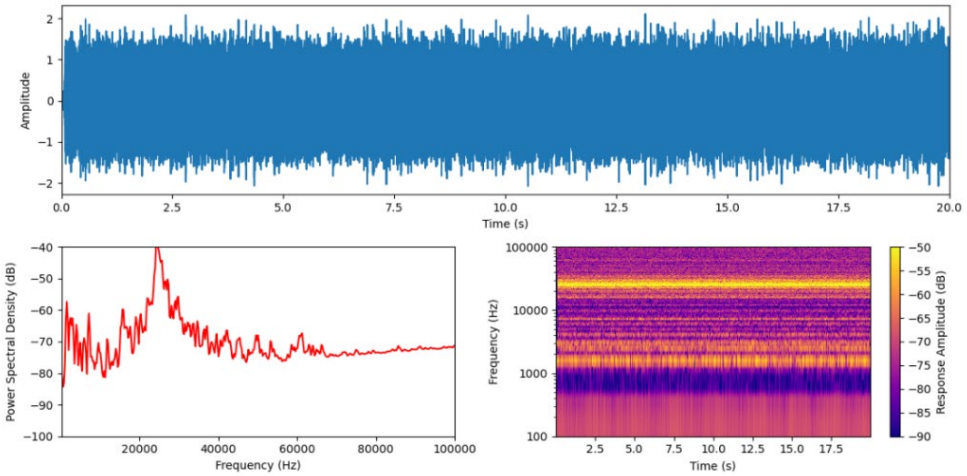

Figure 2. Example of vibrational data collected at point 02 with the plug inserted at 14 cm depth. (top) Vibration time series with amplitude given in voltage. (bottom left) Frequency response spectrum up to 100 kHz. (bottom right) Spectrogram of response amplitude as a function of frequency and time, showing a consistent spectral response over the duration of the measurement. Frequency axis shown in log-scale to better highlight spectral features related to resonance between 1 and 20 kHz.

3.2 Machine Learning Integration

Machine learning provided the analytical power to efficiently extract the relevant SHM features from the large vibrational datasets collected in this study. While graph neural networks (GNNs) had previously been deployed for similar tasks in previously funded MRP work, we chose to augment the model using a Fourier neural operator (FNO) due to its demonstrated ease of implementation and strong predictive performance across a range of physical systems and partial differential equations in recent studies [13–15]. This hybrid physics-informed GNN-FNO framework required minimal backend changes to the workflow and added only small computational expense.

Our GNN-FNO model takes inputs from one or more measured time series, which are then divided into smaller time windows of 4 ms (1000 points) that are processed as labeled training data for initial model fitting, validation data for model tuning, or set aside as unseen test data to evaluate final model performance. Currently, the model can scale up to incorporate at least 100 simultaneous and independent time series measurements. However, we have found certain scenarios where the model can make accurate predictions of experimental conditions using as little as a single time series. Each time series consists of one component of motion at one measurement point and represents one node in the graph. To date, all measurements used with this model consisted of one- or three-component surface velocity measurements made by an LDV. However, the model could seamlessly transition to use strain measurements from distributed acoustic sensors (e.g. fiber optic cables) or data from other types of embedded sensors as long as the inputs consisted of synchronized time series.

The FNO is implemented by first expanding the dimensionality of each input time series in the Fourier (frequency) domain. During the training process, the model learns weights for each Fourier dimension as it identifies the spectral features relevant for the desired predictions. The GNN combines spectral information from each node into node neighborhoods by passing information to each node from all adjacent nodes, retaining the positional information of each node. The node-neighborhood information is passed through a series of convolutional layers, then further combined to make a prediction for the desired property, such as the location of a stress perturbation.

Traditionally, SHM may identify one or more spectral features that indicate stiffening or softening, or are otherwise distorted by stress or damage conditions. The GNN-FNO model identifies many spectral markers across the entire spectrum of the signal *simultaneously* to infer stress and damage conditions, providing a degree of redundancy that minimizes the effects of transient or uncorrelated features on prediction accuracy. We have tested the model by giving it only limited chunks of the expanded dimensionality in the Fourier domain, and the results are clear that markers exist throughout the spectrum.

Prior to analysis of newly collected graphite datasets, improved performance of this ML workflow was confirmed through reprocessing of extensive vibrational data collected on a stainless-steel hexagonal test article in previously funded work [7]. Predictions of positional information—not made available during model training—improved significantly (Figure 3), with correct predictions of location increasing from about 10% of the total with a GNN alone to nearly 90% with the GNN-FNO workflow. Predictions of artificially induced stress levels also improved from 75-85% correct to over 99% correct, even as the model was given a reduced amount of training data from single sensors among the available 105 measurement locations. Additional training details are given for these two models in Appendix A.

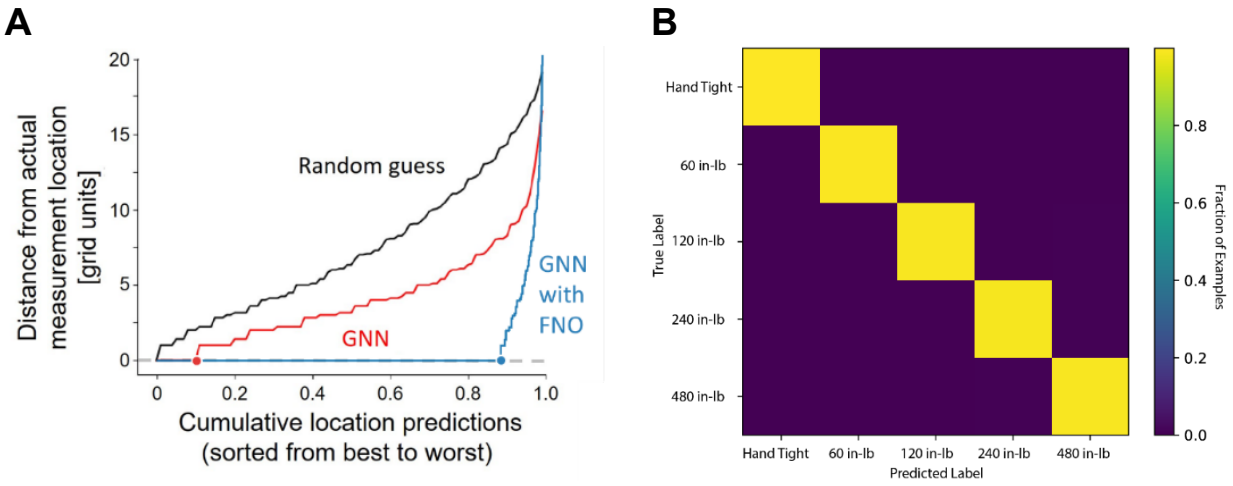

Figure 3. GNN-FNO prediction performance (A) The GNN-FNO model is able to accurately determine the point on the component where the measurement was made. This demonstrates the model learns positional information that is not explicitly presented during the training process. (B) Stress conditions were controlled by applying known levels of torque to a clamping apparatus mounted to the sample. The model has a nearly 100% success rate determining the label associated with the applied torque.

4 Results and Discussion

Baseline acoustic testing of as-machined graphite samples with simple geometries captured linear resonances and modal deflections to determine the feasibility of the experimental design on the more-complex and massive hex block using existing experimental hardware. We found the spherical and rod samples to resonate in a repeatable and detectable manner, suitable for characterization using established resonance-based methods. Results from a subset of spherical samples are given in Figure 4. Elastic characterization using RUS indicated the eight samples had Young's moduli of 10.94±0.14 GPa, closely in line with 10.8 GPa given as a typical value by the supplier. Despite tight clustering of material properties across the samples, part-to-part variations in the vibration responses of the graphite rods and spheres were found to be highly distinctive by the ML model, leading to nearly perfect sample identification predictions (>99% accuracy). We attribute this to minor geometry and density differences

(≤1% across the samples), suggesting sensitivity of the approach to small structural changes. Given the ease of distinguishing the samples using ML-based acoustic characterization, we chose to focus on anomalies introduced within a single sample to remove compounding effects of inter-sample variability.

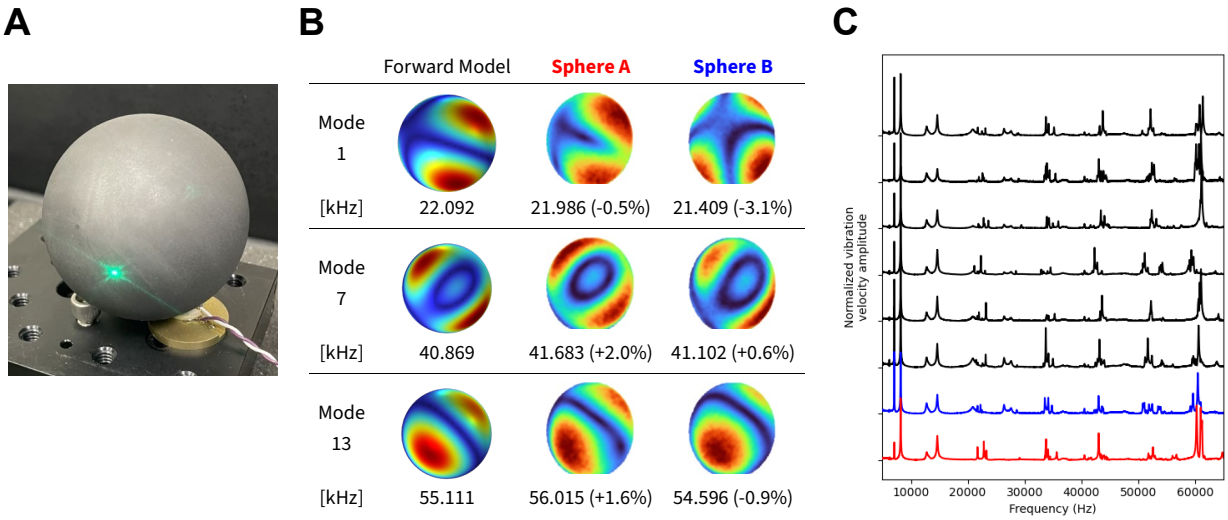

Figure 4. Visual summary of elastic characterization of graphite spheres. (A) Experimental test setup. (B) Selected modeled and measured modes of two samples used for elastic characterization using RUS, confirming the samples retained elastic properties similar to those stated by the supplier after machining. Warm colors in mode shapes represent large vibration amplitudes at the given resonance frequencies. (C) Spectra collected on a set of eight nominally identical samples, with colored curves matching to selected samples from (B). Sharp peaks starting around 22 kHz represent resonances of the samples, while consistent, broad peaks (e.g. near 7 and 12 kHz) originate from the external LDV control mirrors.

Similar baseline modal testing on the hex block confirmed that resonances could be induced throughout the block across a range of input frequencies (Figure 5). While boundary effects from the supporting foam sufficiently shifted frequencies to prevent frequency matching and quantitative material property extraction using RUS, strong mode shape agreement in the audible frequency range confirmed the testing setup was viable and sufficiently repeatable to proceed with ML predictions of anomalous structural conditions within the block.

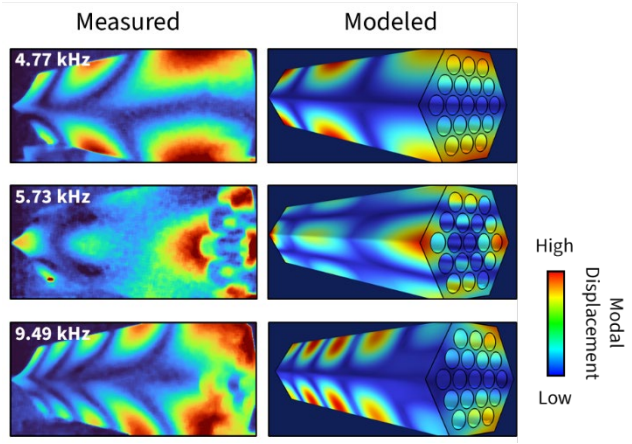

Figure 5. Selected resonance modes measured below 10 kHz (left) and FEA-modeled (right) for the tested hexagonal graphite block. Block geometry is comparable to that described in Trellue et al. [12].

Based partially on measured mode shapes which indicated strong surface responses along the top free edge of the hex block, stress concentrations were introduced within the center, upper hole using the expanding plug assembly. In total, 12 measurement locations were repeated along this top edge for each of the three different plug configurations (not present, 1 cm from end, 14 cm from end) that, due to sample symmetry, represented the 1D spatial extremes of a localized stress anomaly. While this is a simplified test matrix, the resulting model predictions are relevant for both defect identification and localization.

The confusion matrix in Figure 6 summarizes these predictions. Across thousands of examples contained in the test data set from the 12 measured locations, the plug configuration was correctly identified about 75% of the time. The overall presence of the plug was easier to detect from the data, identified at over 90% accuracy. Within the same model run, the plug location was correctly identified at nearly 75% accuracy, suggesting a degree of ambiguity or inconsistency between the sample responses for the two installed plug locations. While this level of accuracy falls far below the 99% accuracy from this same model workflow when applied to artificial stresses within a stainless-steel hex block (Figure 3), a key difference may be the relative volume of the sample that was stressed during each of the experiments. The stainless-steel sample was nearly uniformly compressed, generating stress changes across its entire length, while the hex block experienced anomalous stress only within a small portion of its internal volume.

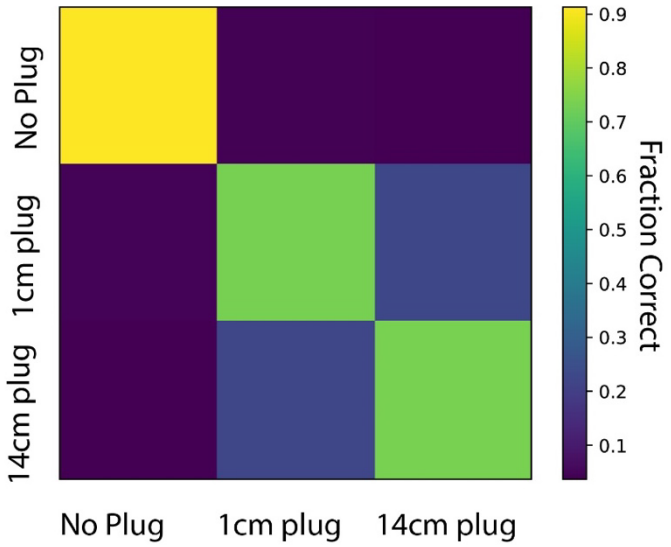

Figure 6. The GNN-FNO model can determine the implemented plug configuration solely from minimal vibration spectra. Predictions are particularly accurate (>90%) for assessing the presence of the plug within the system.

To gain a better understanding of the spectral features learned by the ML model to produce predictions of the plug configuration, we generated a correlation matrix to quantitively compare the similarity of the 36 independently collected spectra (three plug configurations, 12 measurement locations) across the audible frequency range (Figure 7). While a matrix such as this provides an incomplete perspective of the multi-dimensional feature extraction used by the GNN-FNO model, several notable visual patterns are well aligned with the predictive performance of the model.

First, we note the relative dissimilarity of spectral pairs when the plug was absent in at least one of the datasets, with correlations typically falling in the 0.6 – 0.8 range. This contrasts with comparisons of spectral pairs for installed-plug data, where correlations increase into the 0.8 – 0.9 range. As a point of

comparison, a spectrum measured directly on an exposed extension of the plug assembly when it was installed at 14cm produced correlations of 0.2 – 0.3 when compared against data collected on the surface of the block. In Figure 7, outside of the main diagonal showing perfect correlation of each spectrum with itself, we also note more subtle diagonal patterns cutting across the matrix in both directions. This is the result of the combined effects of symmetry in the measurement locations as well as in the resonance mode shapes (see Figure 5). Diagonals going down to the right indicate vibrations measured at a single location remain relatively similar regardless of the plug location, while diagonals going down the left indicate vibration patterns at mirrored locations (e.g. 02 and 09) are relatively similar. Last, as correlation values between and among the two installed plug datasets are very similar, this simple comparison metric using broadband spectral data hints at some of the sources of confusion observed in our model predictions while also highlighting the value in an ML-based approach to signal discrimination and predictive labeling. As opposed to the single frequency range analyzed in Figure 7, the model is capable of simultaneously accessing spectral correlations across many overlapping frequency ranges to identify data features which are most predictive of applied structural change.

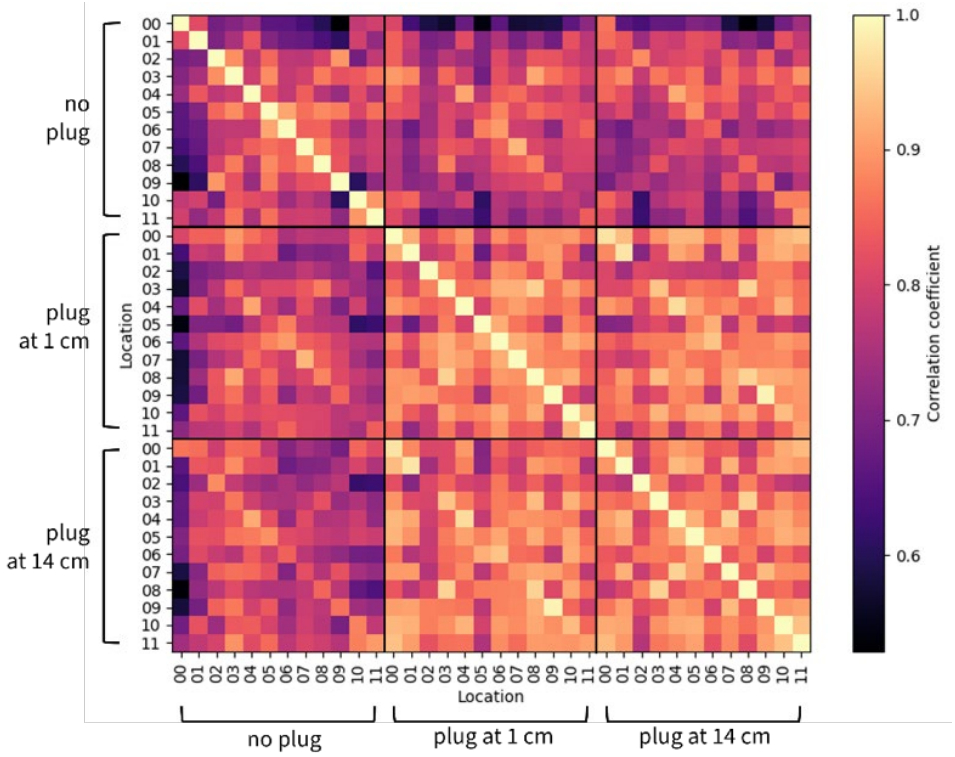

Figure 7. Spectral similarity of vibration measurements made at various points along the surface of the hexagonal block. Lighter colors represent spectral pairs which are more similar, as measured using Pearson correlation coefficients calculated between log-normalized spectra in the audible frequency range from 1 – 20 kHz.

5 Expected Impact

Microreactors may require advanced and innovative automation and significant reductions in staffing relative to commercial reactors to become a competitive energy option despite lower power outputs and limited economies of scale. Critical to this is the development of diagnostic frameworks capable of identifying and localizing structural or system-level defects with high precision and reliability. Embedding distributed sensors and integrating high-frequency acoustic interrogation with complementary sensing modalities will enable early detection of microcracks and other anomalies to increase safety

margins during operation. This and related work previously funded by MRP on enhanced SHM techniques indicate potential benefits to microreactor developers including:

- **Enhanced decision making:** With the ability to leverage recent advances in suitable embedded sensor technologies, such as those in development at ORNL [16,17], signal interrogation techniques, and data reduction methodologies, this framework is suitable for integration with automated control systems to enable continuous monitoring of large-scale components across entire reactor systems. Further, continual detection improvements may be achievable through ongoing model retraining and exposure to greater varieties of operating conditions.

- **Scalability and generalizability:** Incorporation of machine learning algorithms enhances scalability and adaptability across reactor types, establishing a foundation for self-monitoring microreactors capable of sustaining safe and reliable transport and operation under a wide range of conditions. Data processing techniques presented herein are largely sensor agnostic and scalable, meaning there are numerous types of sensor arrays that could be used or combined to enable a predictive model to provide comparable in situ structural information. While extended-duration, high-frequency SHM presents challenges in terms of data storage and computational management, results from this study and previous related work show that the transformation of raw time-series data into frequency-domain representations can be adequate and effective for machine learning–based SHM.

- **Cost reduction:** Greater automation during operation and the ability to provide passive structural evaluations of graphite related to the anomalous stress uniquely induced during transit, startup/shutdown, and operation, reducing need for laborious physical inspections. Outputs may also prove valuable for validation of simulation efforts by microreactor vendors, especially for designs intended be deployed in large numbers.

6 Future Directions

The structural characterization of the graphite hex block presented in this study consisted primarily of highly simplified one-dimensional anomaly location predictions. However, given the geometry of holes within the hex block, localization experiments on repeatable artificial defects can be extended to two and three dimensions to better assess the generalized performance of a GNN-FNO workflow to go from external surface vibration measurements to predictions of the characteristics of an internal structural anomaly. We expect this to provide greater insights into the effects of sensor spatial density, frequency bandwidth of measurements, and the overall generalizability of this approach for anomaly localization.

Increasing the technical viability of this and similar of ML-based SHM approaches should be extended to longer-duration experiments that generate realistic progressive or emergent structural anomalies and enable on-the-fly model refinement and evaluation. This will allow for more robust assessment and validation of these approaches to identify changes from known baseline levels, as would be useful in an automated control system. This may also help in better defining sensitivity to the onset of structural changes during typical startup or operational scenarios. Coupling different sensor types—such as those already incorporated into current reactor design concepts to provide both internal and external data—may further enhance the effectiveness of these methods and enable conclusions to be made between sensor density/layouts and desired defect discrimination and localization targets. SHM efforts within MRP should also integrate with related initiatives, including cross-cutting opportunities in the Structural Materials area of NEAMS, where real-world graphite damage distributions may be useful to validate structural simulations of bulk fatigue and aging.

7 References

[1] Joyce MR, Marrow TJ, Mummery P, Marsden BJ. Observation of microstructure deformation and damage in nuclear graphite. Eng Fract Mech 2008;75:3633–45. https://doi.org/10.1016/j.engfracmech.2007.11.003.

[2] Marsden BJ, Haverty M, Bodel W, Hall GN, Jones AN, Mummery PM, et al. Dimensional change, irradiation creep and thermal/mechanical property changes in nuclear graphite. Int Mater Rev 2016;61:155–82. https://doi.org/10.1080/09506608.2015.1136460.

[3] Metcalfe M. Damage tolerance in the graphite cores of UK power reactors and implications for new build. Nucl Eng Des 2023;406:112237. https://doi.org/10.1016/j.nucengdes.2023.112237.

[4] Moskovic R, Heard PJ, Flewitt PEJ, Wootton MR. Overview of strength, crack propagation and fracture of nuclear reactor moderator graphite. Nucl Eng Des 2013;263:431–42. https://doi.org/10.1016/j.nucengdes.2013.05.011.

[5] Bodel W, Martinuzzi P, Davies B, Steer A, Lowe T, Mummery P. Mimicking irradiation-induced cracking of nuclear graphite using bromine intercalation. Scr Mater 2021;199:10–3. https://doi.org/10.1016/j.scriptamat.2021.113889.

[6] Spencer BW, Munday LB, Singh G, Prithivirajan V. Initial Fracture Propagation Modeling of Graphite Components with Grizzly 2023:INL/RPT-23-74062.

[7] Geimer PR, Delorey AA, Beardslee LB, Ulrich TJ. Demonstration of acoustic monitoring for structural health of microreactors: Through use of neural networks and resonant ultrasound spectroscopy. Los Alamos National Laboratory LA-UR-23-26860: 2023.

[8] General Atomics. Graphite design handbook. DOE-HTGR-88111 1988:1–122.

[9] Black G, Shropshire D, Araújo K, van Heek A. Prospects for Nuclear Microreactors: A Review of the Technology, Economics, and Regulatory Considerations. Nucl Technol 2023;209:S1–20. https://doi.org/10.1080/00295450.2022.2118626.

[10] Arregui-Mena JD, Worth RN, Bodel W, März B, Li W, Campbell AA, et al. Multiscale characterization and comparison of historical and modern nuclear graphite grades. Mater Charact 2022;190. https://doi.org/10.1016/j.matchar.2022.112047.

[11] Pyun D-K, Lee S, Palanisamy RP, Terricabras AJ, van Vuuren JJ, Jacobs D, et al. Comprehensive defect evaluation of advanced nuclear fuels using high-resolution acoustic signals and optimized sensor separation. NDT E Int 2026;157:103502. https://doi.org/10.1016/j.ndteint.2025.103502.

[12] Trellue H, Taylor C, Luther E, Cutler T, Shivprasad A, Jewell JK, et al. Advancements in Yttrium Hydride Moderator Development. Nucl Technol 2023;209:S123–35. https://doi.org/10.1080/00295450.2022.2043088.

[13] Li Z, Kovachki N, Azizzadenesheli K, Liu B, Bhattacharya K, Stuart A, et al. Fourier Neural Operator for Parametric Partial Differential Equations. ICLR 2021 - 9th Int Conf Learn Represent 2021:1–16.

[14] Qin S, Lyu F, Peng W, Geng D, Wang J, Tang X, et al. Toward a Better Understanding of Fourier Neural Operators from a Spectral Perspective. CoRR 2024.

[15] Kaewnuratchadasorn C, Wang J, Kim CW. Neural operator for structural simulation and bridge health monitoring. Comput Civ Infrastruct Eng 2024;39:872–90. https://doi.org/10.1111/mice.13105.

[16] Birri A, Sweeney DE, Hyer HC, Schreiber B, Petrie CM. Development of optical fiber-based sensors for nuclear microreactor structural health monitoring. Oak Ridge National Laboratory ORNL/TM-2023/2984: 2023.

[17] Birri A, Sweeney DC, Hyer HC, Schreiber B, Cakmak E, Petrie CM. A Miniaturized, High-Bandwidth Optical Fiber Fabry-Perot Cavity Vibration Sensor Demonstrated up to 800 °c. IEEE Sens J 2025;25:11082–91. https://doi.org/10.1109/JSEN.2025.3541557.

APPENDIX A

Figure A1 gives two examples of GNN-FNO model fitting, corresponding to the models used to generate the location and stress predictions presented in Figure 3. Appropriate training of the models is, in part, confirmed by the similar evolution in accuracy for the training and validation data sets over the completed training epochs. This is an indication that the model is not overfitting, which would cause it to fail to generalize, where accuracy would decrease significantly if presented with previously unseen validation or test datasets from the same set of experiments.

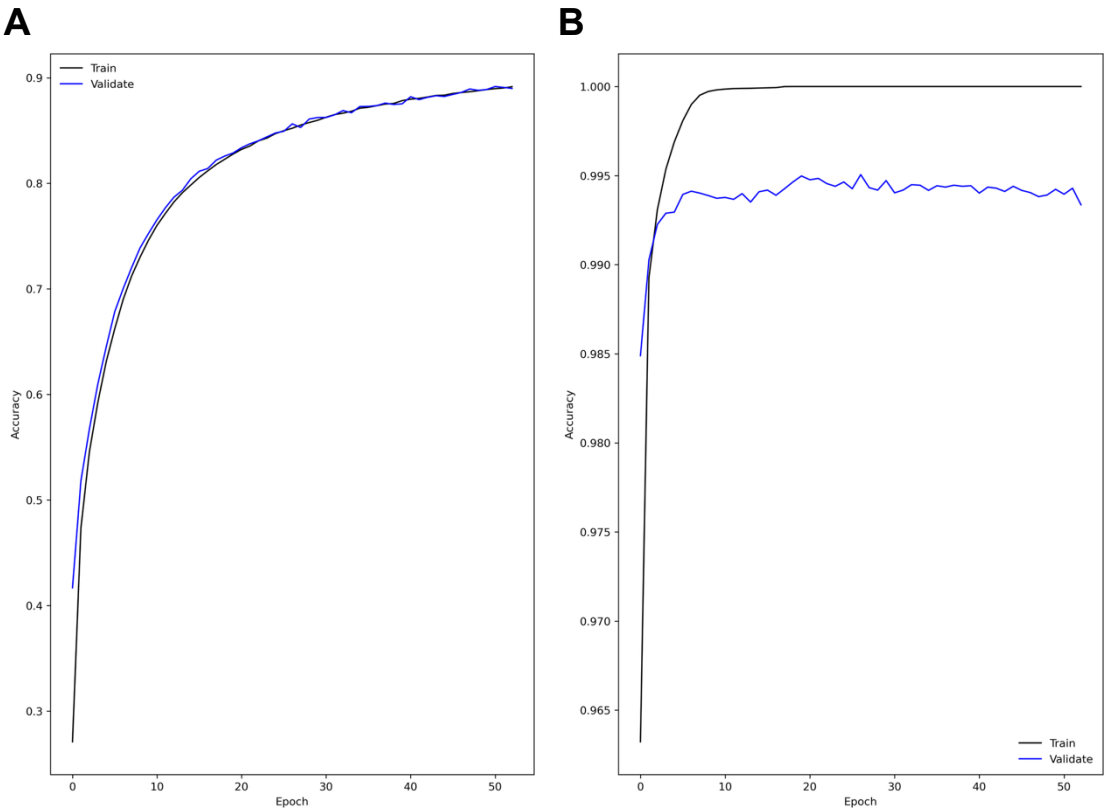

Figure A1. Accuracy of GNN-FNO model predictions as a function of epoch. Prediction accuracy given for location (A) and stress (B) labels from the training (black) and validation (blue) datasets. We note the compressed y-axes in (B) visually exaggerates accuracy differences between training and validation data as compared to (A).