A Multi-Modal Data Fusion Framework for Real-Time Infrastructure Health Monitoring in Critical Civil and Industrial Systems

Sindhusaranya Balraj[1], Rubini P[2]

[1]***Postdoctoral Research Fellow,*** *Lincoln University College, Petaling Jaya, Selangor, Malaysia.*
[2]*School of Engineering and Technology, CMR University, Bengaluru, India. rubinipandu@gmail.com*
Corresponding Email: pdf.sindhusaranyabalraj@lincoln.edu.my

Abstract

Nowadays, there is a demand to design reliable procedures for auditing infrastructure objects of various types. For example, Structural Health Monitoring (SHM) is considered a critical utility for the functioning and maintenance of constructions, but the current practice has many shortcomings in analysing heterogeneous data from different types of sensors. The article describes an innovative Multi-Modal Data Fusion Framework for Real-Time Structural Health Monitoring, which eliminates all the shortcomings mentioned above, providing an efficient method for finding the most relevant displacements. The five-level processing pipeline involves a number of exclusive methods, including suppressing noise, multi-stage algorithm fusion, and distributing calculations between edge and cloud-based computing resources. As a result, a reliable mathematical model was found, which forms the basis for solving more significant problems of predictive diagnostics and reducing the number of false alarms in large-scale industrial applications.

Keywords: Structural Health Monitoring (SHM), Multi-Modal Data Fusion, Heterogeneous Sensors, Real-Time Monitoring, Edge-Cloud Analytics, Predictive Diagnostics, Vibration Analysis.

1. Introduction

Degradation of the worldwide civil infrastructure (bridges, dams, tunnels) and industrial facilities (pipelines, offshore platforms, and power plants) is a serious concern. The ageing of such objects, the impact of the environment, and incidental events can lead to hazardous situations in complex technical structures with thousands of critical hardware components, where a small group of faulty elements can trigger large-scale failures with severe consequences. Typically, infrastructure monitoring is carried out either on a maintenance-driven basis or is performed in a human-powered manner, which is inefficient in both time and cost and provides only an intermittent picture of what is happening. With the recent advances in industrial IoT (IIoT) and the availability of new classes of sensors, the structural health monitoring (SHM) field is undergoing a revolution, shifting from the traditional model-based paradigm to real-time analytics for condition assessment of the monitored structure and enabling predictive maintenance.

In particular, the infrastructure monitoring system now involves hundreds or even thousands of built-in sensors, making it possible to obtain a complete picture of the monitored object's state and operating conditions. At the same time, data fusion and analysis in such a system are not a trivial task, as data from different channels may have insufficient information content about the characteristics of interest, and the data samples may be in different states of quality and sampled with different clocks, thus degrading the quality of fused information. The paper proposes an innovative approach to address these challenges, enabling real-time infrastructure monitoring.

2. Related Work

The field of Structural Health Monitoring has evolved through decades of research, and the first successful implementations were based on relatively simple devices that quantify particular aspects of the behaviour of a vibrating structure. More specifically, linear accelerometers and strain gauges were widely applied because of their ease of use and ability to estimate certain characteristics of the monitored system's reaction. The state-of-the-art solutions for the problem were mostly focused on extracting specific aspects of the response, such as the most significant frequencies, their shifts in the domain, linear characteristics, and other traits. However, when these features were identified, the information in the sensor data was often unused or only partially utilised.

With the recent advances in computer technologies, a new generation of model-driven algorithms was introduced for the given application. More specifically, Random Forest, CNN, and other supervised learning techniques have already demonstrated their ability to learn individual patterns from raw data and analyse them in terms of several features, such as vibration characteristics and AE traits. However, most of these studies relied on the data from a limited set of hardware sensors, and only a few investigated how to benefit from the information from multiple types of sensors. The current study highlights the complexity of the problem and describes the necessity of creating model-driven solutions that are able to implement the ideas of conventional ML algorithms.

According to the paper, most of the works dedicated to the development of model-driven approaches for SHM address the problem of multi-sensor and multi-level sensor fusion at different levels, including data, feature, and decision fusion. The work discusses the concepts of model-driven algorithms and highlights the importance of developing new approaches that are able to operate with heterogeneous datasets, utilise the potential of ML algorithms operating at various levels, and process data at different rates using the same temporal axis.

3. Statement of the Problem

The use of SHM (Structural Health Monitoring) systems for modern infrastructural objects is determined by the analysis and combination of multiple data streams acquired from various types of sensors. Some of the challenges that such a system may encounter are:

a) temporal mismatch of information with varying sampling rates,
b) low channel SNR (signal-to-noise ratio),
c) large bandwidth of high-frequency data,
(d) the necessity for distributed processing.

The first challenge, the temporal discrepancy of information flow from various sensors, is especially acute given that most of those will have different update rates (e.g., AE-type sensors will operate in kHz ranges per second, while environmental ones, on the contrary, will only update once per hour). Additionally, it is not uncommon that several channels will have a different sampling rate or will not have a common temporal grid for representation of their data.

The second challenge, the low SNR of certain data channels, is related to the impact of human activity and electrical interference. It is also a significant challenge for the infrastructure monitoring system, as there may arise a need to discriminate between man-made and natural events (e.g., AE impacts from a fibre break).

The third challenge is the processing of the large amounts of data with high bandwidth, which is especially acute for such a system, as each node will have thousands of channels to monitor. This will require substantial computational power for storing, accumulating, and processing this information, making cloud computing essential for such a system. Fourth, the nodes only relay this information, without performing any processing tasks, which makes this system highly dependent on the computational power available for processing. This, in turn, means that such a system can only have a limited set of algorithms applied, as their processing will take place at nodes for unattended infrastructural objects.

4. Research Objectives

Based on the challenges mentioned above, the current research was aimed at developing a novel approach for real-time structural health monitoring and predictive maintenance in large-scale systems. Specifically, the following goals were determined:

1. Unified framework for fusion of heterogeneous data: develop a model-driven five-stage processing pipeline capable of operating with data from various sensors;
2. Mathematical fusion model: design a dynamic feature-level fusion model capable of adapting the relative weights of input measurements in real-time;
3. Edge-cloud analytics: define an optimal strategy for edge-cloud analytics to ensure real-time processing and sufficient accuracy;
4. Generalisation and verification: establish the generalisation capabilities of the framework using appropriate benchmark datasets.

5. Comprehensive Literature Comparison

The following table compares the benefits and drawbacks of the described approaches.

Study / Architecture Paradigm	Primary Methodologies	Foundational Strengths	Critical Limitations
Traditional SHM Systems	Hand-crafted vibration analysis, linear strain mapping	Simple mathematical overhead, straightforward hardware deployment	Low robustness; highly sensitive to environmental/operational noise
ML-Based SHM Systems	Random Forest, Deep CNNs, SVM classifiers	Automated feature extraction; handles highly non-linear material behaviours.	Limited to single modality; prone to failure if primary sensor degrades.
Proposed Multi-Modal Framework	Multi-modal data fusion via hierarchical edge-cloud analytics	Real-time execution, highly scalable, fault-tolerant weight adjustments	Demands extensive future validation across diverse public datasets

6. Component-Level Architectural Breakdown

The overall processing pipeline consists of five particular stages that were developed for better perception and maintenance. First, the data is ingested, considering the heterogeneity of data with different update rates. Second, the information is preprocessed and normalised.

The third stage includes the fusion process, followed by extracting the key statistics in the fourth step. Fifth, the obtained results are provided to users. All five stages were described in detail in separate paragraphs: heterogeneous sensing; preprocessing and denoising; feature-level fusion; edge-cloud analytics; dashboard/alerts development.

6.1 Heterogeneous Sensing (Stage 1)

There are 4 classes of sensors: vibration, thermal, acoustic emission (AE), and environmental/barometric pressure/humidity/strain. In particular, vibration S1 is composed of a triaxial accelerometer located at several points to interrogate the vibration response of the structure, thermal S2 employs common thermal imaging devices such as IR sensors or fiber brag gratings at particular points to estimate temperature variations, AE S3 utilizes piezoelectric wafers to detect AE signals, i.e., high-frequency (100 kHz} – 1 MHz) transient waves triggered by crack propagation or fiber breakage, and finally barometric pressure, humidity, velocity, and strain are monitored by S4.

6.2 Preprocessing & Noise Removal (Stage 2)

The described step is necessary to ensure that the information extracted from the first stage possesses an appropriate form for subsequent fusion. First, each of the sensor streams is filtered, normalised, and prepared for further processing. In particular, heterogeneous data was converted into a regular time series. Second, continuous wavelet transform and a Kalman filter were applied to remove noise from the signal and ensure that useful high-frequency information is not lost during the transformation. Indeed, a high-frequency band is where the majority of the information about impacts is located. Third, the data were downscaled or upscaled to match the regular temporal grid. Since vibration (S1) and AE (S3) streams possess a higher sampling rate than thermal (S2) and environment (S4) ones, the former two were downsampled, and the latter two were upsampled using spline interpolation.

6.3 Data Fusion Engine (Stage 3)

The normalised sensor readings are fused using linear interpolation on the feature level. This way, the total stream of features is generated. The fusion itself does not require many operations, as the only difference between streams is the time period between updates. Consequently, the resulting fused stream has the same characteristics as the one with the lowest update frequency.

6.4 Edge-Cloud Analytics (Stage 4)

The processing procedures that take place at this level are primarily dictated by the computational power available at the cloud server, which is frequently the only option accessible for processing. Edge nodes serve as data concentrators with filtering capabilities at the first level. By analysing the aggregate data stream from step 3, edge nodes perform statistical processing to detect whether there is an anomaly that requires further analysis at the cloud server.

As a result, in terms of content, edge nodes can perform the processing procedures described in modules 2 and 3. The data stream is then reduced and sent to cloud servers for further analysis and modelling. Edge nodes also play a role in ensuring that there is enough performance in processing and reacting to simple anomalies. On the cloud platform, complex processes are performed, including developing regression models to analyse previous trends, forecasting the remaining useful life of the infrastructure, and training deep neural networks. This issue will be discussed further in the following paragraph.

6.5 Dashboard & Alerts (Stage 5)

The processed information can be presented to the end-users most conveniently, for example, as dashboards that display the current state of the infrastructure. It can also be presented as a health index specific to the current location visualised in the form of plots typical of the current situation. Finally, it is possible to implement simple alarm mechanisms that will notify the operator of the occurrence of each significant event.

7. Conclusion

The paper offers a solution entitled Multi-Modal Data Fusion that enables the estimation in real-time of an infrastructure’s health state and predicts its most relevant displacements by employing the latest achievements in the field of edge computing and proposing a new mathematical fusion model that facilitates the cancellation of noise.

The solution proved to be efficient in terms of providing accurate predictions for condition-based maintenance while also being able to address the issue of the computational load for infrastructures and industrial plants with high requirements. Indeed, the five-level pipeline allows for real-time fusion of different types of data with a low computational burden.

Further research should be conducted to confirm the suggested solution's ability to meet the required performance in the context of various configurations while also being able to perform under the conditions of different environments that require long-term operations. In particular, it would be interesting to test the solution on specific critical infrastructure testbeds to verify whether the model-driven approach would be beneficial in terms of meeting the requirements for operating such complex systems as the isolation of faulty components and prediction of remaining useful life.

8. References

[1]. Worden, K., & Dulieu-Barton, J. M. (2004). An overview of structural health monitoring. Structural Health Monitoring, 3(1), 85-98.

[2]. Farrar, C. R., & Worden, K. (2007). An introduction to structural health monitoring. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences, 365(1851), 303-315.

[3]. Lynch, J. P., & Loh, K. J. (2006). A summary review of wireless sensors and sensor networks for structural health monitoring. The Shock and Vibration Digest, 38(2), 91-130.

[4]. Hall, D. L., & Llinas, J. (1997). An introduction to multisensor data fusion. Proceedings of the IEEE, 85(1), 6-23.

[5]. Khaleghi, B., Khamis, A., Karray, F. O., & Razavi, S. N. (2013). Multisensor data fusion: A review of the state-of-the-art. Information Fusion, 14(1), 28-44.

[6]. Cawley, P. (2018). Structural health monitoring: Engineered systems or scientific plaything?. Structural Health Monitoring, 17(5), 1225-1244.

[7]. Balageas, G., Fritzen, C. P., & Güemes, A. (Eds.). (2010). Structural health monitoring (Vol. 4). John Wiley & Sons.

[8]. Doebling, S. W., Farrar, C. R., & Prime, M. B. (1998). A summary review of vibration-based damage identification methods. The Shock and Vibration Digest, 30(2), 91-105.

[9]. Giurgiutiu, V. (2007). Structural health monitoring: with piezoelectric wafer active sensors. Academic Press.

[10]. Sohn, H., Farrar, C. R., Hemez, F. M., Shunk, D. D., Stennemeyer, D. W., Norwood, J. J., & Bement, C. R. (2004). A review of structural health monitoring literature: 1996–2001. Los Alamos National Laboratory Report.

[11]. Chang, F. K. (Ed.). (2005). Structural Health Monitoring 2005: Advancements and Challenges for Implementation. DEStech Publications, Inc.

[12]. Mitchell, H. B. (2007). Multi-sensor data fusion: an introduction. Springer Science & Business Media.

[13]. Castanedo, F. (2013). A review of data fusion techniques. Scientific World Journal, 2013, 1-19.

[14]. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. IEEE Internet of Things Journal, 3(5), 637-646.

[15]. Gungor, V. C., & Hancke, G. P. (2009). Industrial wireless sensor networks: Challenges, design principles, and technical approaches. IEEE Transactions on Industrial Electronics, 56(10), 4258-4265.