A Privacy-Preserving Framework Using Federated Learning for Structural Health Monitoring with Miter Gates Application

YICHAO YANG, HAIWEI LIN, GUOFENG QIAN, ZHEN HU and MICHAEL D. TODD

ABSTRACT

Structural Health Monitoring (SHM) systems play a critical role in maintaining the safety and operational efficiency of infrastructure such as miter gates managed by the United States Army Corps of Engineers (USACE). These gates, integral to navigation and flood control systems, demand robust and efficient monitoring techniques to predict failures and thus optimize maintenance schedules. Any unexpected shutdown costs nearly three million dollars per day to the US economy. The existing traditional data centralized machine learning approaches for SHM often face challenges, including data privacy concerns, high communication costs, and computational limitations. This study mainly explores the application of Federated Learning (FL) techniques to SHM systems using three finite element miter gate models for the loss-of-contact damage detection problem. This approach addresses these challenges by enabling decentralized training of machine learning models across multiple assets. FL ensures data privacy by keeping sensitive information local while leveraging shared global models through aggregation methods. Our results demonstrate that FL can achieve comparable prediction accuracy to centralized methods while maintaining data privacy and reducing communication overhead. This framework improves model accuracy by incorporating diverse data distributions from different gates and their operational conditions. The study also highlights the scalability of FL in handling large-scale SHM systems, making it a viable solution for USACE to extend the lifespan of their critical assets.

Yichao Yang, Haiwei Lin, Guofeng Qian, Michael D. Todd, Department of Structural Engineering, University of California at San Diego, 9500 Gilman Drive, La Jolla, CA 92093. Zhen Hu, Department of Industrial and Manufacturing Systems Engineering, University of Michigan-Dearborn, 4901 Evergreen Rd, Dearborn, MI 48128. Corresponding author email: mdtodd@ucsd.edu

1. INTRODUCTION

The United States Army Corps of Engineers (USACE) operates and maintains critical waterway infrastructure, including dams, locks, and especially miter gates, which are essential for navigation and flood control [1–3]. Many of these structures are aging and well beyond their original design life (e.g. some even exceeding 80 years in service). [4]. Unexpected failures or unscheduled maintenance events can result in significant economic losses, estimated at nearly three million dollars per day of downtime [1,5,6]. The challenges faced by USACE include the increasing frequency of maintenance needs, as well as escalating costs associated with inspections and repairs [7, 8]. An efficient and reliable structural health monitoring system has become indispensable.

One of the most commonly observed damage mechanisms is the development of gaps due to loss of contact at the quoin block [6, 9]. The quoin block serves as the contact interface between the gate and the supporting lock wall, and any degradation at this location can severely compromise the gate's structural integrity [10]. Additionally, hydrostatic loadings, resulting from differences in upstream and downstream water levels, impose significant stresses on the gates. Structural Health Monitoring (SHM) systems have emerged as life-cycle management tools for continuously assessing the integrity of large-scale infrastructure such as miter gates [1, 3, 11]. SHM enables early damage detection and thus reduces the risk of catastrophic failures.

However, traditional centralized SHM approaches still face major challenges such as concerns over data privacy, high communication costs, and lack of adaptability to localized degradation [8]. Federated Learning (FL) offers a promising solution by enabling decentralized model training directly on local devices. This method can preserve individual data privacy and reduce communication overhead [12, 13]. Specifically, FL has the potential to revolutionize SHM by combining insights from distributed data without aggregating raw sensor measurements. At the same time, it can also allow robust generalization across diverse operational conditions. In this paper, we present a privacy-preserving framework for SHM based on Federated Learning, specifically tailored for the miter gate damage detection problem. We used our high-resolution nonlinear Abaqus finite element model simulate three digital miter gate models, each subjected to different material property variations to mimic real-world diversity [14]. Damage scenarios are created by introducing varying gap lengths at the quoin block, and corresponding strain data are collected from selected gauges [2]. Aggregation technique such as FedAvg is employed to synthesize knowledge from distributed models without violating data privacy [15, 16].

Our preliminary results demonstrate that the federated learning approach achieves comparable, and sometimes superior, prediction performance relative to traditional centralized models while maintaining strict data privacy. This work highlights the significant potential of federated learning to address the unique challenges associated with population-based SHM in large-scale waterway infrastructures.

2. FEDERATED LEARNING

Federated Learning (FL) is a decentralized machine learning framework that enables multiple clients to collaboratively train a shared global model without transferring their local data to a central server [17, 18]. In traditional centralized learning, all data must be collected and aggregated in a central location before training, which raises concerns related to data privacy, communication costs, and regulatory compliance. FL addresses these challenges by keeping data localized at the client side, exchanging only model parameters or updates during the training process [19].

In an FL system, each client, such as a mobile device (or a structural system) independently trains a local model using its own private dataset [12, 20]. Periodically, the local model parameters are sent to a central server, where they are aggregated (e.g., using Federated Averaging, FedAvg) to form a new global model. These updated global model parameters are then redistributed to clients, and the process continues iteratively until convergence. By decoupling model training from direct access to raw data, FL significantly enhances privacy protection and thus reduces network bandwidth requirements. In addition, FL also enables the deployment of machine learning models in environments with sensitive or distributed data [13].

FL techniques are critical for enabling large-scale and privacy-preserving AI systems in real-world scenarios. Real-life applications of FL have grown rapidly across several sectors. In healthcare, FL enables collaborative medical research across hospitals without sharing sensitive patient data and provides diagnostic recommendations [21]. In mobile devices, FL powers next-word prediction and personalized services (e.g., Google GBoard) without uploading user text inputs to central servers [20]. In finance, FL is employed for fraud detection across multiple banks without exposing proprietary customer information [22].

In the context of SHM, FL can be a transformative technology. Large-scale civil infrastructures are spatially distributed and often monitored individually. Implementing FL in SHM allows each structure to autonomously contribute to the development of robust predictive models without sharing raw structural response data [23, 24]. Then, the facilitate collaborative learning is produced across diverse structural assets.

3. FEDERATED LEARNING FOR SHM OF MITER GATES

This application of FL to SHM of miter gates presents a novel approach to address the inherent challenges of miter gates' diverse structural conditions. In traditional SHM systems, sensor data collected from individual structures must be transmitted to a central server for analysis and model training. This centralized approach is often impractical for critical infrastructure such as miter gates, where data privacy and bandwidth limitations are major concerns [25].

By integrating FL into SHM, each miter gate independently trains a local predictive model using its own sensor data—such as strain and hydrostatic load levels—without transmitting raw measurements to a central server. Instead, only the locally updated model parameters are communicated [26]. A central aggregator collects these parameters and performs an aggregation (e.g., via Federated Averaging, FedAvg) to produce an improved global model. The updated global model is then redistributed back to each gate for the next local training round. This iterative process continues until model convergence. Figure 1 illustrates four key steps of the FL framework for SHM of miter gates. Each gate acts as an individual local client, performing local model training based on its local sensor measurements (i.e. strain data). Through secure model aggregation, the framework builds a powerful predictive model capable of estimating damage characteristics such as gap length at the quoin block, while maintaining strict privacy guarantees.

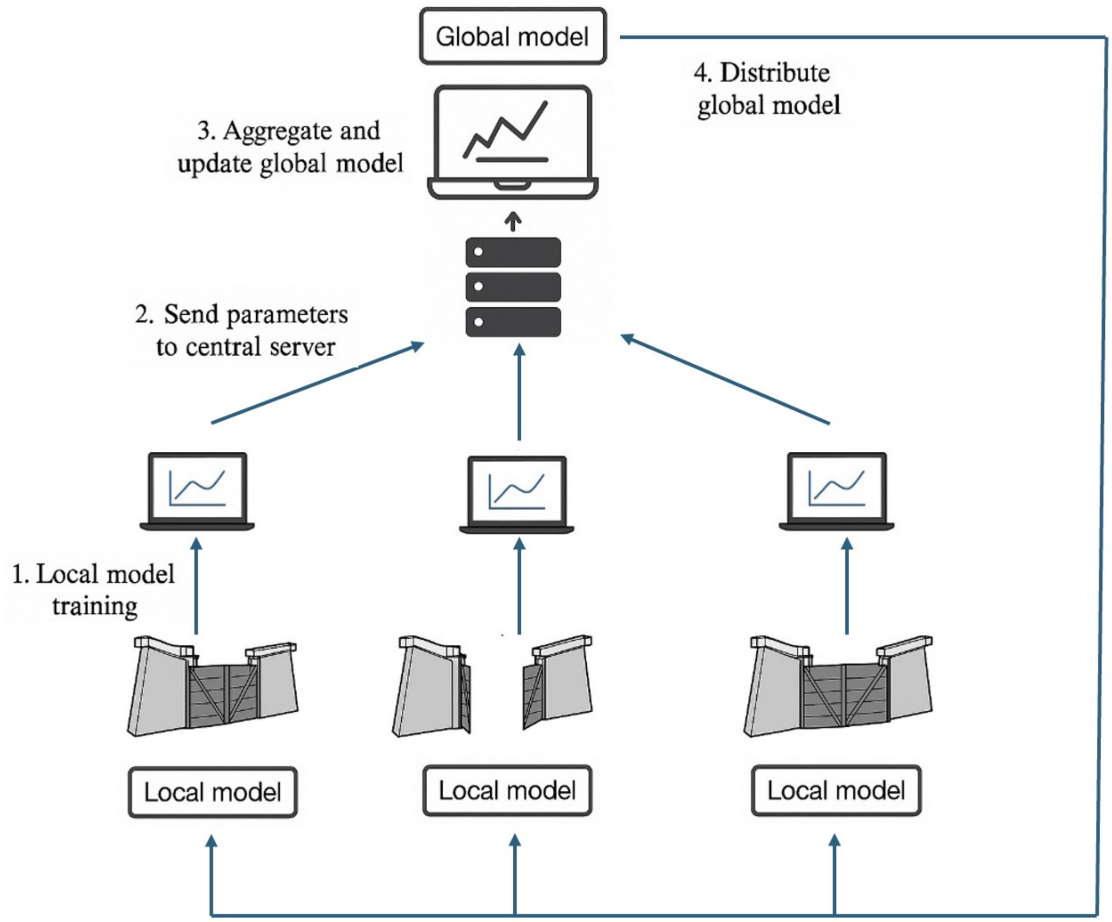

Figure 1. Four key steps of the proposed FL framework for SHM of miter gates.

Applying FL to miter gates enables several critical advantages.

- Data Privacy Preservation: Sensitive structural data remains local, minimizing exposure risks.
- Communication Efficiency: Only model weights are shared, significantly reducing data transmission volumes compared to raw sensor data transfer.
- Model Robustness: By training on diverse but related datasets across different gates, the global model can generalize better to various structural conditions such as degradation levels and loading environments.

3.1. Aggregation Method: Federated Averaging

The Federated Averaging (FedAvg) algorithm is the core aggregation method used to combine local models into a single global model during federated training [15,16]. After each communication round, every client (miter gate) returns an updated local model with parameters θ_i, trained using its own local dataset D_i.

Let $n_i = |D_i|$ be the number of samples in client i, and let $n = \sum_{i=1}^{N} n_i$ be the total number of samples across all clients. The FedAvg method aggregates the local weights θ_i into a global model θ_G as a weighted average as follows:

$$\theta_G = \sum_{i=1}^{N} \frac{n_i}{n} \cdot \theta_i. \tag{1}$$

This aggregation is performed component-wise across all weights of the neural network. Mathematically, if each θ_i is a vector of parameters, the global parameter θ_G is also a vector of the same size, computed as:

$$\theta_G^{(j)} = \sum_{i=1}^{N} \frac{n_i}{n} \cdot \theta_i^{(j)} \quad \text{for each layer/parameter } j. \tag{2}$$

One of the main benefits of FedAvg is handling non-IID and unbalanced data distributions [27]. It also reduces communication by allowing multiple local epochs between aggregations.

4. MITER GATES FOR CASE STUDY

To demonstrate the proposed FL-based SHM framework for miter gates, we created three similar but different miter gates by using the ”Greenup” miter gate model. The Greenup miter gate is a miter gate located on the Ohio river. Figure 2 shows a physics-based finite element (FE) model of the miter gate developed in Abaqus. The structure is primarily composed of steel, so linear elastic material properties are used to simulate the mechanical behavior of the steel. Additionally, as the structure is mainly made up of welded steel plates, the model is constructed using 8-node shell elements (S8R). Identical loads and damage scenarios were applied to each model as input variables, while the resulting strain values at the strain gage locations were used as output variables.

Damage is introduced into the model through the presence of a gap at the quoin block, representing the loss of contact between the gate and the supporting wall. (shown in Figure 3). This gap can significantly affect stress distributions and overall system stability. Hydrostatic loads are applied independently at the upstream (H_{up}) and downstream (H_{down}) sides of the gate to simulate water pressure differentials encountered during lock operations [6](shown in Figure 4). The primary outputs of the simulation are the strain responses recorded at 42 strategically placed strain gauges across the structure, as shown in Figure 5. To replicate actual strain gage measurements, the FE model was locally partitioned at sensor locations, and material orientations were assigned so that the E_{11} strain component aligned with the gage’s sensing direction.

4.1. Dataset Generation Using the Greenup Model

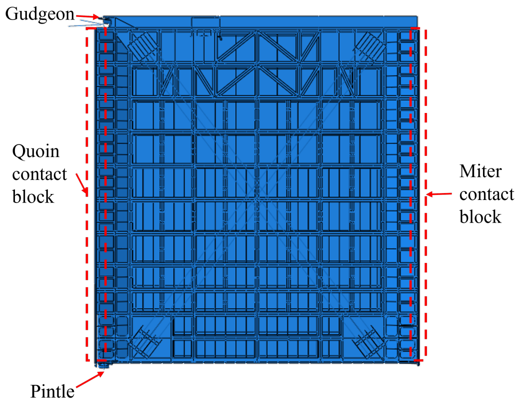

Figure 2. Schematic and main components for Greenup model

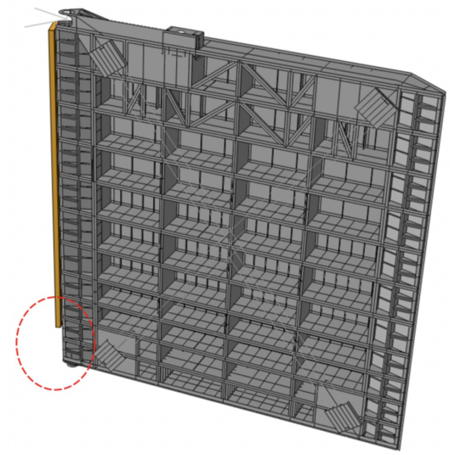

Figure 3. Gap damage of miter gate [28]

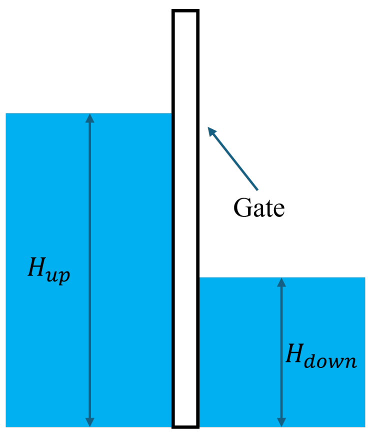

Figure 4. Hydrostatic loads on gate

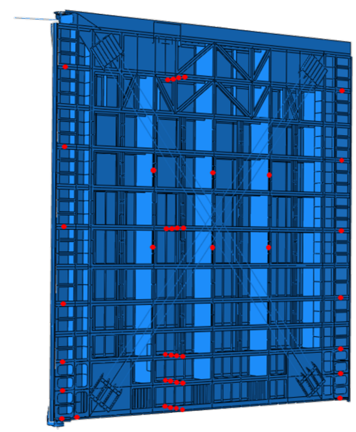

Figure 5. Strain Gauges in Greenup model

Using the Greenup FE model, synthetic datasets were generated by varying three key input parameters:

- Upstream water height (H_{up}): 500 to 598.43 inches.
- Downstream water height (H_{down}): 98.43 to 228.35 inches.
- Gap ratio (g): 0 to 0.2375, corresponding to a physical gap length between 0 and 180.975 inches (computed as 0.2375×762 inches, where the length of the entire contact block is 762 inches).

A total of 256 input samples were generated using Sobol sequence sampling to ensure a uniform and comprehensive coverage of the input space [29,30]. For each sampled input, the Greenup model outputs the corresponding strain component E_{11} from the 42 sensors.

4.2. Simulation of Multiple Miter Gates via Material Variation

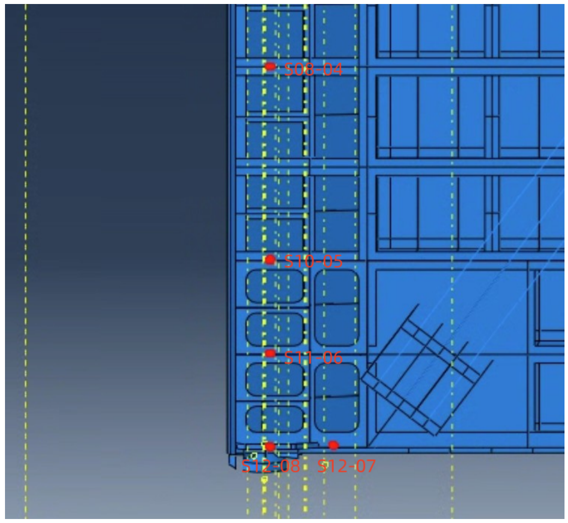

Figure 6. Selected sensors near damage (gap) area.

To simulate operational variability across different gates, the Greenup model was further modified by adjusting the Young's modulus of the material. Three different cases were created:

- Case E1: Young's modulus = 29,000 ksi (reference).
- Case E0.8: Young's modulus = 23,200 ksi (80% of reference).
- Case E0.6: Young's modulus = 17,400 ksi (60% of reference).

For each material property scenario, a corresponding dataset of 256 samples was generated under the same Sobol sampling scheme. This modification represents different aging and degradation states that miter gates may experience in practice.

Thus, three independent but related datasets were obtained, each simulating the behavior of a different miter gate under a range of loading and damage conditions.

5. DATASET PROCESSING AND FL MODEL TRAINING

5.1. Sensor Selection and Input-Output Structure

After generating the datasets from the Greenup Abaqus forward model simulations, the next step was to formulate an inverse problem to predict the gap length based on selected sensor measurements and hydrostatic loading conditions. Five strain gauges located near the damage area were selected, specifically sensors "S12-08," "S12-07," "S11-06," "S10-05," and "S08-04" (locations shown in Figure 6).

For each miter gate model, the input dataset was structured as:

- Inputs: [H_{up}, H_{down}, S12-08, S12-07, S11-06, S10-05, S08-04]
- Output: Gap length (computed as gap ratio $\times$ 762 inches)

5.2. Dataset Segmentation for Federated Learning Demonstration

Our selection strategy is motivated by real-world miter gate operations. In practice, different gates across the USACE inventory are subjected to varying degrees of degradation and loading histories based on age, material deterioration, and environmental factors. Some gates experience early-stage small gaps due to minor wear, while others, especially older or more heavily used gates, present mid-to-large size gaps near failure. Therefore, each local dataset realistically represents different operational damage severities, reflecting the actual diversity seen across deployed gates (as shown in Figure 7).

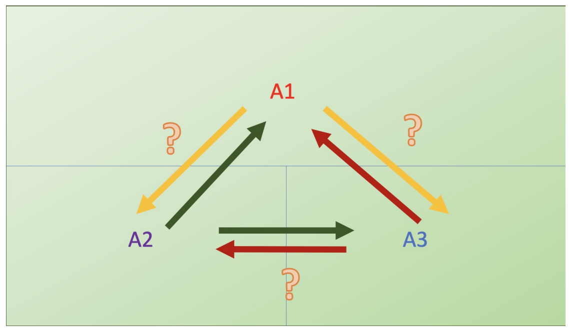

Figure 7. Select data from different damage range.

To emphasize the benefits of federated learning in dealing with data heterogeneity, the datasets were partitioned based on different gap length ranges (listed in Table I):

TABLE I. Selected Training and Test Datasets

Dataset	Gap Length Range (inches)	Training Samples	Testing Samples
E1 (X1, y1)	0–90	128	256
E0.8 (X2, y2)	60–150	127	256
E0.6 (X3, y3)	120–180	84	256

By simulating these different segments separately for each ”local” model, the federated learning framework can better capture and aggregate knowledge from non-identically distributed (non-IID) data sources, which is a key challenge in practical SHM applications. We select these specific sample counts for training sets is based on the number of available samples within the chosen gap length ranges (shown in Table I). For example, E1 has 128 matching samples in betweeen gap range (0–90) inches.

Thus, the training sets use all available samples within their respective ranges, maximizing the data utilized for local models. For testing, however, we evaluated the trained models against the entire dataset of 256 original samples from each gate (not just the range-restricted samples). The training and test samples are shown in Figure 8. In this way, the training datasets from three miter gates are from three different regions with minor overlaps. Besides, it simulates real-world scenarios where a deployed model may encounter data outside its initial training regime, thus testing robustness.

This split strategy reflects a practical federated learning condition where local clients (gates) specialize in certain operational conditions but the global server aims for broader generalization across all possible conditions.

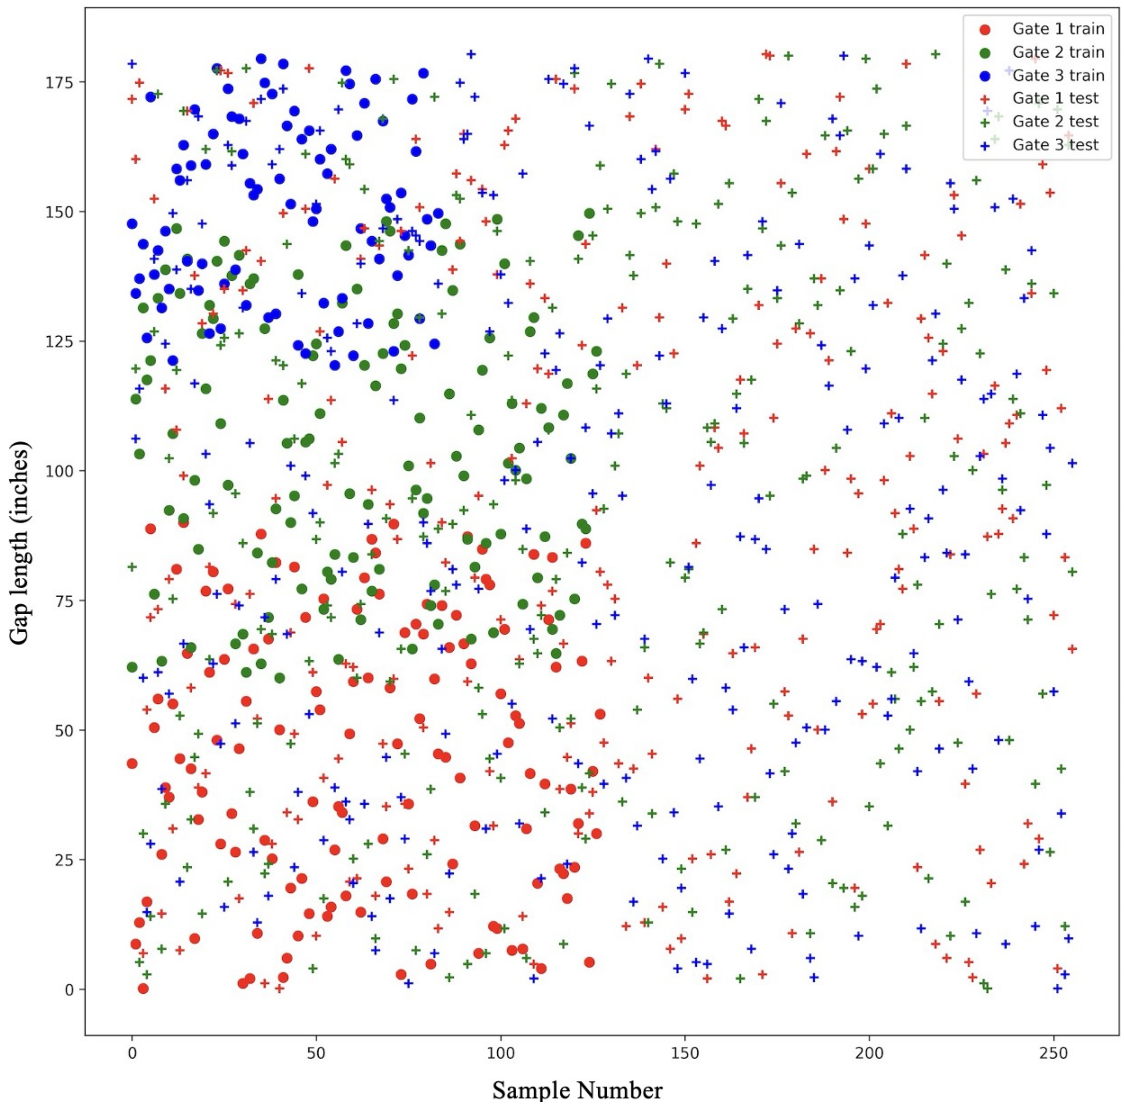

Figure 8. Training and Test data samples

5.3. ML model architecture

Given the significant differences in magnitude between different input features, normalization was applied to ensure numerical stability and effective model training. The hydrostatic pressure inputs (Hup and Hdown) are large-valued, typically ranging from 100 to 600 inches, while the strain gauge outputs are extremely small, often measured in the order of microstrain. Without proper normalization, models tend to favor features with larger numerical scales, leading to biased training with slow convergence [31]. Standard normalization is performed to the data before the training of the machine learning models. For FL in this case study, each local client employs a feedforward multi-layer perceptron (MLP) model for the regression task of predicting miter gate gap lengths [32–34]. The architecture is relatively simple yet effective for capturing the underlying relationships between hydrostatic loads, strain values, and gap lengths.

Figure 9 shows a simple model architecture for this work. The input layer consists of seven features: upstream and downstream water heights, and strain measurements from five selected gauges near the damage zone. This input is passed through two hidden layers: the first hidden layer has 64 neurons with ReLU activation, and the second hidden layer has 32 neurons, also using ReLU activation. To enhance generalization and prevent overfitting, a dropout layer with a 0.2 dropout rate is applied after the first hidden layer during training. The final output layer consists of a single neuron with a linear activation function, representing the predicted gap length. The model is trained to minimize the mean squared error (MSE) loss function, using the Adam optimizer with a learning rate of 0.01.

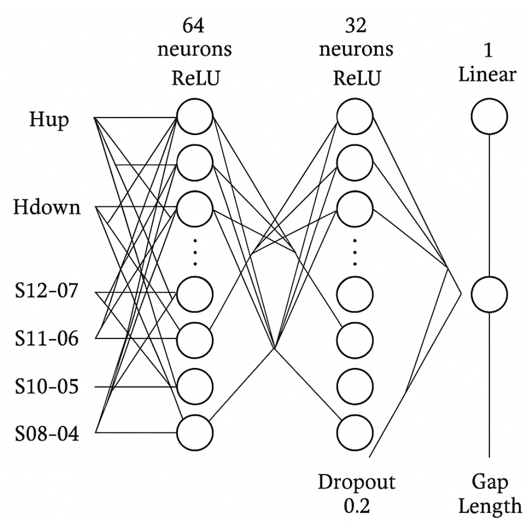

Figure 9. A simple MLP architecture

Overall, this lightweight MLP design balances model complexity and computational efficiency, making it well-suited for distributed training across multiple miter gates with limited local data. The prediction task then involves feeding new test inputs $\mathbf{x}_{\text{test}} \in \mathbb{R}^d$ into the trained MLP model to estimate the gap length $\hat{y}_{\text{test}}$.

5.4. Summary of Workflow

One of the key parts in this work is to develop the full training and communication procedure of our privacy-preserving FL framework that is customized for SHM applications in miter gates. The goal is to collaboratively train a global model across multiple local datasets without sharing raw data, thus we share the local model parameters. Figure 10 illustrates the general procedure of the FL framework, which illustrates the interaction between the central server and local clients (miter gates).

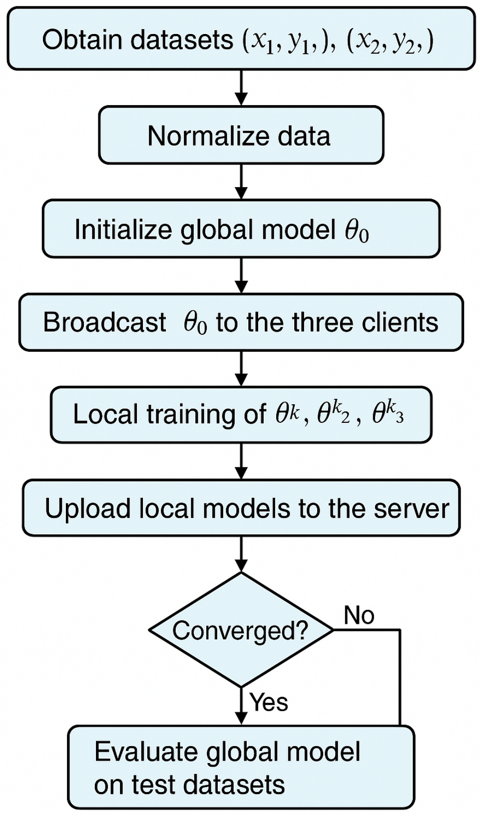

Figure 10. A workflow of FL framework for SHM on miter gates.

The entire workflow proceeds as follows:

1. Central server initializes a global model with random parameters θ_0.
2. In each iteration round t, the server broadcasts the current global model parameters θ_t to all clients.
3. Each client updates the model parameters locally using its dataset D_i, trains for several epochs, and sends the updated weights θ_t^i back to the central server.
4. The server aggregates the weights (e.g., via FedAvg) to produce an updated global model parameters θ_{t+1}.
5. This process continues until convergence.
6. Distribute the global model to each local clients (i.e. miter gates), and then perform local fine-tuning

6. RESULTS

We explored and compared three learning strategies between local, centralized, and federated methods. Each local model was trained using only the training subset of its respective training dataset (E1, E0.8, or E0.6). The training was performed using the same MLP architecture described previously with early stopping to prevent overfitting. In contrast, a centralized model was trained using the combined training data from all three datasets. After training, the centralized model was fine-tuned individually on each training set to simulate local domain adaptation. In addition, in the federated setup, local models were trained independently using their local data and their model parameters were synchronized through federated averaging (FedAvg) over multiple communication rounds. Once convergence was detected, the final global model parameters were distributed back to the local models for fine-tuning. Each locally fine-tuned model was then evaluated on its corresponding test set.

The prediction performance of the models is summarized in Table II. It clearly illustrates comparisons of prediction performance between local, centralized, and federated models across the three local dataset predictions. Performance is evaluated in terms of mean squared error (MSE) on the respective test sets.

TABLE II. Comparison of Prediction Performance (MSE) Across Learning Strategies

Model Type	**Training Data**	**Test Data**	**MSE**
Local Model 1	$X_1^{\text{train}}, y_1^{\text{train}}$	$X_1^{\text{test}}, y_1^{\text{test}}$	8.0100
Local Model 2	$X_2^{\text{train}}, y_2^{\text{train}}$	$X_2^{\text{test}}, y_2^{\text{test}}$	0.1155
Local Model 3	$X_3^{\text{train}}, y_3^{\text{train}}$	$X_3^{\text{test}}, y_3^{\text{test}}$	1.2381
Model 1 (centralized)	$(X_{all}^{\text{train}}, y_{all}^{\text{train}})$ + fine-tuned $(X_1^{\text{train}}, y_1^{\text{train}})$	$X_1^{\text{test}}, y_1^{\text{test}}$	0.4695
Model 2 (centralized)	$(X_{all}^{\text{train}}, y_{all}^{\text{train}})$ + fine-tuned $(X_2^{\text{train}}, y_2^{\text{train}})$	$X_2^{\text{test}}, y_2^{\text{test}}$	0.1081
Model 3 (centralized)	$(X_{all}^{\text{train}}, y_{all}^{\text{train}})$ + fine-tuned $(X_3^{\text{train}}, y_3^{\text{train}})$	$X_3^{\text{test}}, y_3^{\text{test}}$	0.1824
Model 1 (federated)	FL + fine-tuned $(X_1^{\text{train}}, y_1^{\text{train}})$	$X_1^{\text{test}}, y_1^{\text{test}}$	0.2364
Model 2 (federated)	FL + fine-tuned $(X_2^{\text{train}}, y_2^{\text{train}})$	$X_2^{\text{test}}, y_2^{\text{test}}$	0.1050
Model 3 (federated)	FL + fine-tuned $(X_3^{\text{train}}, y_3^{\text{train}})$	$X_3^{\text{test}}, y_3^{\text{test}}$	0.1853

Figure 11 below presents a 3-by-3 grid of subfigures, each showing the true test values (x-axis) versus predicted values (y-axis) of gap lengths. From the first row to the third row, it compares plots for local, centralized, and federated methods, respectively. From the first column to the third column, it represents each method performances to the local ”miter gate” number 1 to 3. In each subplots, the red dashed line represents the ”ideal” predictions, while the blue dots indicate the actual predictions.

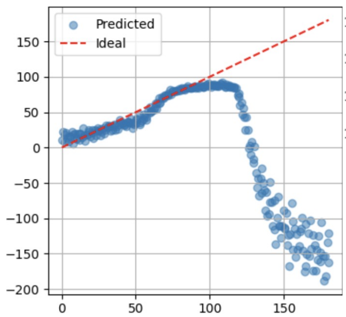

(a) Local Model 1
(MSE = 8.0100)

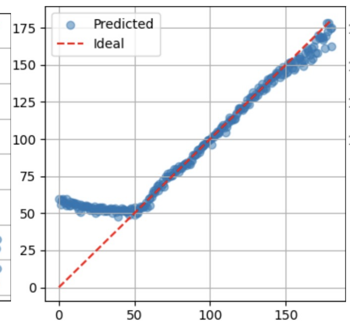

(b) Local Model 2
(MSE = 0.1155)

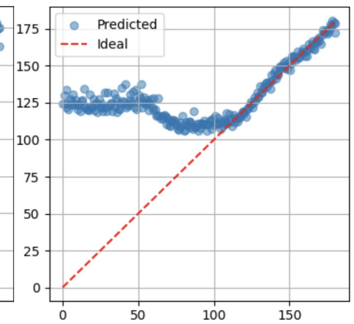

(c) Local Model 3
(MSE = 1.2381)

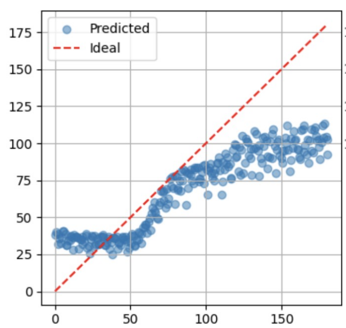

(d) Model 1 (Centralized)
(MSE = 0.4695)

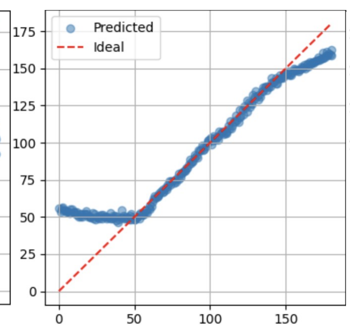

(e) Model 2 (Centralized)
(MSE = 0.1081)

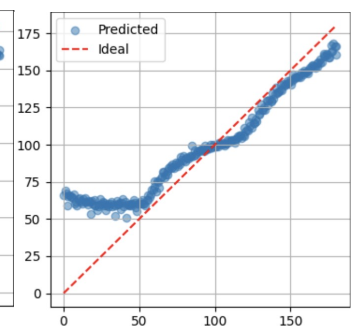

(f) Model 3 (Centralized)
(MSE = 0.1824)

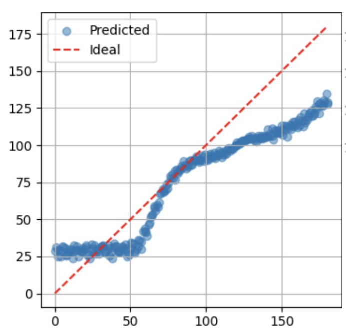

(g) Model 1 (Federated)
(MSE = 0.2364)

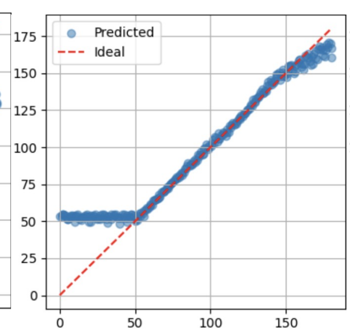

(h) Model 2 (Federated)
(MSE = 0.1050)

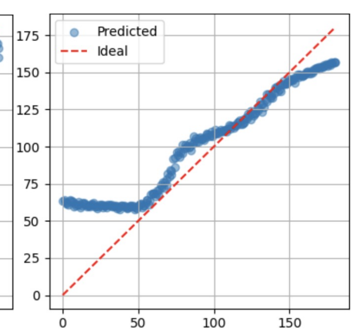

(i) Model 3 (Federated)
(MSE = 0.1853)

Figure 11. True vs. predicted gap lengths for Local, Centralized, and Federated methods across all local miter gates.

6.1. Results discussion

- *Local Models:* Local models trained on isolated datasets show relatively poor generalization performance, particularly for Model 1 and Model 3. Their high MSE values reflect overfitting to the local domain, which is a common issue when data diversity is limited.

- *Centralized Models:* Centralized training using all available data yields better generalization with significantly lower MSE across all three test datasets. However, this approach compromises data privacy and requires data transfer.

- *Federated Models:* The federated learning models show performance comparable to centralized models, with only marginal increases in MSE. Especially when comparing the Model 1 prediction using different methods (first column), the FL approach achieves much better performance than the other two methods. This suggests that FL effectively leverages diverse data distributions.

These results demonstrate the effectiveness of the FL method in achieving near-centralized accuracy (some even better) while maintaining decentralized data and preserving privacy. It also concludes that the FL methods outperform local model approaches in distributed data scenarios.

7. CONCLUDING REMARKS

In this work, we proposed and demonstrated a privacy-preserving federated learning framework for damage estimation in miter gate structures. We also presented a detailed implementation of the FL workflow. Three datasets were created by modifying the Young's modulus of the physically based Greenup model to reflect different structural behaviors. We trained and evaluated local models, centralized models, and federated models, comparing their performance using prediction accuracy. Results showed that the FL model achieved prediction performance comparable to the centralized model, while maintaining data privacy and enabling decentralized training.

In the next stage of this research, we plan to introduce further heterogeneity into the datasets by modifying additional physical parameters beyond Young's modulus. This will make the three miter gate models more distinct and realistic. Furthermore, we will expand the simulation environment by incorporating three entirely different Abaqus finite element models, each representing a unique miter gate. This extension will allow us to evaluate the scalability and adaptability of federated learning in broader and more practical SHM scenarios across diverse infrastructure assets.

ACKNOWLEDGMENT

This work was supported by the US Army Engineer Research and Development Center (ERDC) under cooperative agreement W9132T-22-20014. Z. H. was also partially supported by the U.S. National Science Foundation under Grant CMMI-2423521. The supports are gratefully acknowledged.

REFERENCES

1. Eick, B. A., Z. R. Treece, B. F. S. Jr, M. D. Smith, S. C. Sweeney, Q. G. Alexander, and S. D. Foltz. 2018. "Automated damage detection in miter gates of navigation locks," *Structural Control and Health Monitoring*, 25(1):e2053.
2. Yang, Y., M. Chadha, Z. Hu, and M. D. Todd. 2023. "An optimal sensor design framework accounting for sensor reliability over the structural life cycle," *Mechanical Systems and Signal Processing*, 202:110673.
3. Parno, M., D. O'Connor, and M. Smith. 2018. "Monitoring the structural health of lock gates with Gaussian processes and model reduction," *ASCE Conference Proceedings*.
4. Dang, T. V., Q. A. Mai, P. G. Morato, and P. Rigo. 2020. "Updating the Reliability of Aging Miter Gates in the Presence of Corrosion and Fatigue," in *ICSCEA 2019: Proceedings of the International Conference on Sustainable Civil Engineering and Architecture*, Springer Singapore, pp. 1247–1253.
5. Schwieterman, J. P., S. Field, L. Fischer, and A. Pizzano. 2010. "An analysis of the economic effects of terminating operations at the Chicago river controlling works and O'Brien locks on the Chicago area waterway system," *Chaddick Institute for Metropolitan Development, DePaul University*, 24:2010.
6. Yang, Y., M. Chadha, Z. Hu, M. A. Vega, M. D. Parno, and M. D. Todd. 2021. "A probabilistic optimal sensor design approach for structural health monitoring using risk-weighted f-divergence," *Mechanical Systems and Signal Processing*, 161:107920.
7. Vega, M. A., Z. Hu, Y. Yang, M. Chadha, and M. D. Todd. 2022. "Diagnosis, prognosis, and maintenance decision making for civil infrastructure: Bayesian data analytics and machine learning," in *Structural Health Monitoring Based on Data Science Techniques*, Springer, pp. 45–73.
8. Estes, A. C., D. M. Frangopol, and S. D. Foltz. 2004. "Updating reliability of steel miter gates on locks and dams using visual inspection results," *Engineering Structures*, 26(3):319–333.
9. Eick, B. A., Z. R. Treece, B. F. S. Jr, M. D. Smith, S. C. Sweeney, Q. G. Alexander, and S. D. Foltz. 2018. "Automated damage detection in miter gates of navigation locks," *Structural Control and Health Monitoring*, 25(1):e2053.
10. Greimann, L. F., J. H. Stecker, A. M. Kao, and K. L. Rens. 1991. "Inspection and rating of miter lock gates," *Journal of Performance of Constructed Facilities*, 5(4):226–238.
11. Yang, Y., M. Chadha, Z. Hu, and M. D. Todd. 2022. "An optimal sensor placement design framework for structural health monitoring using Bayes risk," *Mechanical Systems and Signal Processing*, 168:108618.
12. Li, L., Y. Fan, M. Tse, and K.-Y. Lin. 2020. "A review of applications in federated learning," *Computers & Industrial Engineering*, 149:106854.
13. Lu, H., A. Thelen, O. Fink, C. Hu, and S. Laflamme. 2024. "Federated learning with uncertainty-based client clustering for fleet-wide fault diagnosis," *Mechanical Systems and Signal Processing*, 210:111068.
14. Corporation, D. S. S. 2011. *Abaqus 6.11*.
15. Wang, H., M. Yurochkin, Y. Sun, D. Papailiopoulos, and Y. Khazaeni. 2020. "Federated learning with matched averaging," *arXiv preprint arXiv:2002.06440*.
16. Wen, J., Z. Zhang, Y. Lan, Z. Cui, J. Cai, and W. Zhang. 2023. "A survey on federated learning: challenges and applications," *International Journal of Machine Learning and Cybernetics*, 14(2):513–535.
17. Mammen, P. M. 2021. "Federated learning: Opportunities and challenges," *arXiv preprint arXiv:2101.05428*.

18. Zhang, C., Y. Xie, H. Bai, B. Yu, W. Li, and Y. Gao. 2021. "A survey on federated learning," *Knowledge-Based Systems*, 216:106775.
19. Li, T., A. K. Sahu, A. Talwalkar, and V. Smith. 2020. "Federated Learning: Challenges, Methods, and Future Directions," *IEEE Signal Processing Magazine*, 37(3):50–60.
20. Yang, T., G. Andrew, H. Eichner, H. Sun, W. Li, N. Kong, D. Ramage, and F. Beaufays. 2018. "Applied Federated Learning: Improving Google Keyboard Query Suggestions," *arXiv preprint arXiv:1812.02903*.
21. Xu, J., B. S. Glicksberg, C. Su, P. Walker, J. Bian, and F. Wang. 2021. "Federated learning for healthcare informatics," *Journal of Healthcare Informatics Research*, 5:1–19.
22. Yang, W., Y. Zhang, K. Ye, L. Li, and C.-Z. Xu. 2019. "FFD: A Federated Learning Based Method for Credit Card Fraud Detection," in *Big Data–BigData 2019: 8th International Congress, Held as Part of the Services Conference Federation, SCF 2019, San Diego, CA, USA, June 25–30, 2019, Proceedings 8*, Springer International Publishing, pp. 18–32.
23. Sharma, A., S. K. Singh, A. Chhabra, S. Kumar, V. Arya, and M. Moslehpour. 2023. "A Novel Deep Federated Learning-Based Model to Enhance Privacy in Critical Infrastructure Systems," *International Journal of Software Science and Computational Intelligence (IJSSCI)*, 15(1):1–23.
24. Novikova, E., E. Doynikova, and S. Golubev. 2022. "Federated Learning for Intrusion Detection in the Critical Infrastructures: Vertically Partitioned Data Use Case," *Algorithms*, 15(4):104.
25. Anaissi, A., B. Suleiman, and M. Naji. 2021. "Intelligent Structural Damage Detection: A Federated Learning Approach," in *Advances in Intelligent Data Analysis XIX: 19th International Symposium on Intelligent Data Analysis, IDA 2021, Porto, Portugal, April 26–28, 2021, Proceedings*, Springer International Publishing, vol. 19, pp. 155–170.
26. Anaissi, A., B. Suleiman, and W. Alyassine. 2023. "Personalised federated learning framework for damage detection in structural health monitoring," *Journal of Civil Structural Health Monitoring*, 13(2):295–308.
27. Zhao, Y., M. Li, L. Lai, N. Suda, D. Civin, and V. Chandra. 2018. "Federated Learning with Non-IID Data," *arXiv preprint arXiv:1806.00582*.
28. Vega, M. A. 2020. *Diagnosis, prognosis, and maintenance decision making for civil infrastructure*, Ph.D. thesis, University of California, San Diego.
29. Burhenne, S., D. Jacob, and G. P. Henze. 2011. "Sampling based on Sobol' sequences for Monte Carlo techniques applied to building simulations," in *Building Simulation 2011*, IBPSA, vol. 12, pp. 1816–1823.
30. Dige, N. and U. Diwekar. 2018. "Efficient sampling algorithm for large-scale optimization under uncertainty problems," *Computers & Chemical Engineering*, 115:431–454.
31. Singh, D. and B. Singh. 2020. "Investigating the impact of data normalization on classification performance," *Applied Soft Computing*, 97:105524.
32. Hoskere, V., B. Eick, B. F. S. Jr, M. D. Smith, and S. D. Foltz. 2020. "Deep Bayesian neural networks for damage quantification in miter gates of navigation locks," *Structural Health Monitoring*, 19(5):1391–1420.
33. Pinkus, A. 1999. "Approximation theory of the MLP model in neural networks," *Acta Numerica*, 8:143–195.
34. Ghate, V. N. and S. V. Dudul. 2010. "Optimal MLP neural network classifier for fault detection of three phase induction motor," *Expert Systems with Applications*, 37(4):3468–3481.