Damage localization using differentiable physics and displacement-based structural identification

Borna Rahnamay Farnod [a] [iD], Wesley F. Reinhart [b,c] [iD], Rebecca K. Napolitano [a] [iD],*

[a] *Department of Architectural Engineering, The Pennsylvania State University, University Park, 16802, PA, USA*
[b] *Department of Materials Science and Engineering, The Pennsylvania State University, University Park, 16802, PA, USA*
[c] *Institute for Computational and Data Sciences, The Pennsylvania State University, University Park, 16802, PA, USA*

ARTICLE INFO

ABSTRACT

As the built environment ages and incurs natural wear and tear, it becomes imperative to employ measures that allow cost-effective, rapid, and automated inspection methods to ensure human safety. Displacement-based damage localization methods can meet these needs through non-contact monitoring techniques such as digital image correlation that provide high-resolution displacement measurement across large-scale civil structures rapidly and with great accuracy. In this work, we explore the use of displacement information as the input to an inverse solver that reconstructs the spatially varying material properties of the target system. We demonstrate the usefulness of the displacement-based data and our inverse solver on a collection of simple case studies: a fixed-fixed beam, single-story, single-span frame, and multi-story, multi-span frame. We systematically study the sensitivity of the system to damage and its location, the ability of the model to identify and localize multiple simultaneous damage, the damage severity, and the robustness of the optimizer with respect to initialization. We also study challenges likely to face the method in real-world applications such as noisy measurements, reconstruction without full-field data, extension to different system sizes and geometries, and sensitivity of the algorithm to mismatch in resolution. Investigation of the single-span frame shows the solver can achieve localization accuracy of 83.8% when there is no measurement noise and 71.3% with 2% noise. Additionally, the proposed method achieves 82.9% accuracy using surface displacement-only which is slightly lower than 84.2% localization accuracy when the full-depth displacement is available. Results demonstrate the effectiveness of the proposed solver in identifying subsurface damages with high accuracy using surface displacement only.

1. Introduction

The U.S. is facing a growing problem with its declining infrastructure [1]. Analytical models, backed by advancements in data collection and processing techniques, provide a unified means for the damage prognosis of structures [2]. Additionally, they can be used to store, visualize, and infer key properties of structures needed for structural health monitoring (SHM) purposes, which in turn combats the risk surrounding declining infrastructure.

Damage diagnosis in analytical models involves working backward: deducing model details from specific measurements. A common damage diagnosis method is to model the system using the finite element method (FEM) and update the model's parameters, which are representative of the damage status. Additionally, accurate modeling of the structure' details determine how well it captures the actual structure's behavior, hence providing accurate information about the existing and future state of the system [2].

Updating the FE model parameter, also called parameter estimation, can be done using dynamic and static data [3]. However, dynamic data often comes with disadvantages [4]. A major problem with dynamic data is the increased error as the result of approximating the damping matrix as part of the dynamic equation [5]. Another issue with dynamic data is that some damage information is only accessible at higher modes, which are often expensive or difficult to induce on a structure [6]. In contrast, static parameter estimation requires updating stiffness and mass matrices, which are easier to recover due to fewer unknowns [7].

Parameter estimation using static data consists of three stages [3]: identifying the objective function, dealing with incomplete or sparse measurements, and choosing a scheme to solve the optimization problem. Two common objective functions to select from are a displacement and a force error estimator [8]. In this work, we choose the displace-

---

ment metric to measure the discrepancy between the actual and analytical models' responses. Moreover, studies show that measuring displacement errors is preferable, especially in the presence of noise [3,7].

Early research in displacement-based damage localization focused on using sparse static measurements to recover stiffness parameters [9,10]. Significant methodological advances included sensitivity-based approaches for preserving matrix properties [7] and gradient-based optimization formulation [8]. While these developments enabled parameter identification even with limited observations and noisy measurements, the fundamental limitations of sparse measurements prevented precise element-level damage localization and quantification [11,12]. Even with various mathematical enhancements, displacement-based methods typically require prior accurate knowledge of structural stiffness [13].

Further progress in parameter identification using static measurements was seen with advances in close-range photogrammetry. Close-range photogrammetry captures the spatial variation of targets of interest [13] and provides accurate full-field displacement measurements. Digital Image Correlation (DIC) systems provide superior data compared to other optical full-field measurement technologies [14]; it additionally has a measurement noise often within the order of 0.01 to 0.05 pixels [15,16] which can be reliably used as the ground truth for updating FEM's response [16]. As a result, DIC became a promising method for FE model updating and parameter estimation [11,14,17, 18]. DIC measurements were particularly useful in studying deformation, cracks [19–21] and damage in beams [22,23], bridges [24,25], and wind turbine blades [26].

Research in the field of full-field displacement measurement for updating structural models has progressed from simple to more complex systems, but there remains a gap in the literature regarding its application to FEMU of systems with a large number of parameters [12,27] with a focus on using deterministic optimization [28]. An early example is the work of Zaletelj et al. [12], which used full-field measurements to recover modal shapes of a simply supported beam with 999 elements. That study particularly focused on evaluating the effectiveness of their methodology in recovering Young's Modulus of damaged DoFs. Dizaji et al. used full-field displacement response to update the full model of an I-shaped steel beam subjected to flexural loading [29]. That study updated a model with 4300 elements, demonstrating the scalability of the method in terms of the number of parameters [30]. In further exploration, the same authors [30] investigated the application of 3DIC (Three-Dimensional Digital Image Correlation) for identifying subsurface damage [11]. They employed a topology optimization formulation for internal damage identification of a laboratory-scale T-bone element, validating their approach with a FE model with up to 28390 elements. Their study also marked DIC as a promising technology for non-destructive evaluation (NDE) of civil structures. We examined the potential of DIC-based NDE methods using synthetic data. That study explored recovering homogeneous and inhomogeneous material properties in a model beam with 800 elements [31].

Most of these studies have primarily studied small-scale structural systems, such as beams. The adaption of 3DIC technology for large-scale civil structures, while promising, also faces non-technical challenges. While this method enables non-contact distant measurement of structures with minimal interruption to infrastructure and users, several practical limitations exist [32]. The initial setup requires expertise in system calibration, including installation of targets or application of speckle patterns onto structures for measurement purposes [33]. Environmental factors pose significant challenges — the method demands small deformations and a consistent field of view [32]. Weather variations, lighting changes, localized reflection, glare, and poor speckle patterns can all contribute to measurement inaccuracies [11].

When operating in real-world conditions, accurately identifying and isolating mechanisms becomes complex. These mechanisms encompass both system loads (such as pedestrian and vehicular traffic) and ambient factors like wind effect on both the structure and measurements [33]. Furthermore, while this approach does not require specialized equipment like GPR and other NDE tools, it demands additional labor in other areas. Specifically, it necessitates extensive preprocessing, postprocessing, and computational work, requiring expertise in both FE modeling and optimization techniques [11].

This paper focuses on the local FEMU of frame structures using full-field displacement data. We aim to extend the utility of our proposed framework [31] for parameter identification in large civil structures. Several key questions are addressed in this regard. First, we investigate whether surface deformation information alone can detect internal damage in a frame structure. Secondly, we provide a comparative analysis of full-field measurement with discrete measurement. Finally, we conduct a series of systematic numerical experiments to assess the impact of noise in measurement and FE modeling. Additionally, we perform sensitivity analyses to evaluate various aspects of the framework and its parameters.

It is worth emphasizing that our proposed approach is distinct and more comprehensive than the work done by Dizaji et al. [11, 34]. Damage localization and quantification using topology optimization assumes that material property is related to the material density through the SIMP (Solid Isotropic Material with Penalization) topology optimization formulation. We lift this limiting assumption by enabling the direct estimation of material property through displacement data. The following section will present a detailed overview of our proposed method, including a brief introduction to differentiable physics-based models and their application to parameter identification.

2. Methodology

We use an automated method to calculate the sensitivity of model parameters with respect to the output, assuming a linear elastic structure. This assumption is valid for structural systems with small deformations. Consequently, we employ the linear elasticity formulation within the differentiable physics solver to represent the structure's behavior. We rely on `dolfin-adjoint` to calculate the system's sensitivity using the adjoint method, which is computationally efficient for computing gradients of systems with many unknowns.

The computational cost of the adjoint method primarily increases with the number of outputs variables, while it is largely independent of the number of input variables, i.e., the number of elements. Although adjoint method is highly scalable in this sense, one must be mindful of its memory requirements and the increased computational cost of the forward problem. In contrast, computation cost of methods like finite difference directly scales with the number of input variables. This characteristic makes the adjoint method well-suited for the optimization of systems with a large number of elements [35,36].

In this study, we aim to recover the spatially varying inhomogeneous material properties throughout the structure using local FEMU. For complete details of the computational framework, refer to Ref. [31]. Here, we will summarize the Partial Differential Equation (PDE) of the system, the gradient-based optimization, and the regularization method used in this work. We will also discuss the evaluation criteria for quantifying the solution's fitness.

2.1. Differentiable physics model

We use an automated method to perform gradient-based updates on the differentiable physics model. The `dolfin-adjoint` [37] software provides a differentiable interface for models in the FEniCS [38] finite element solver. Initially, it is only necessary to consider the physics model. We implement linear elasticity within FEniCS, using the principle of virtual work to derive the equations of strain energy, as defined in our earlier work [39]:

Find $u \in V$

$$s.t. \ a(u,v) = \int_{\Omega} \sigma(u) : \epsilon(v) \, d\Omega = \int_{\Omega} f \cdot v \, d\Omega \quad \forall v \in \hat{V},$$
$$\sigma(u) = \lambda \, \text{tr}(\epsilon(v)) \, I + 2\mu\epsilon(v),$$
$$\epsilon(v) = \frac{1}{2}(\nabla v + (\nabla v)^T).$$

(1)

In the above equation, the goal is to find u, the displacement vector field, within the solution space V. This field u must satisfy the equilibrium of forces, represented by the bilinear form $a(u, v)$ for internal virtual work. The stress tensor $\sigma(u)$ and strain tensor $\epsilon(v)$ are defined using the deformation gradient ∇u and Lamé's parameters λ and μ. I is the identity matrix, f represents the body force per unit volume, and Ω is the problem's domain.

These equations may include spatially varying material properties, specifically elastic modulus $E(x)$, where x indicates spatial coordinates within Ω. For simplicity, we introduce a dimensionless parameter $\theta(x)$, termed fractional elastic modulus, defined as $\theta(x) = E(x)/E_0$, where $E(x)$ is the elastic modulus at position x and E_0 is some reference elastic modulus. Representing the spatially varying material properties in terms of θ allows for consideration of any materials, given a differentiable physical model, using a coefficient factor. In practice, we represent the body as a collection of discrete elements, such that $\theta(x)$ is implemented as a vector θ_i of continuous values between $[0, 1]$, with one entry for each element i in the mesh.

We can explicitly introduce θ in Eq. (1) as a coefficient as follows:

$$\sigma(u, x) = \theta(x) \left[\lambda_0(\mathrm{tr})\epsilon(v)I + 2\mu_0\epsilon(v) \right], \tag{2}$$

The right-hand side of the above equation directly modifies the elasticity tensor C based on θ.

2.2. Gradient-based optimization

We compute the gradient within the PDE solver using the `dolfin-adjoint` library. Gradient descent is performed using the Adam optimizer [40], a common gradient-based optimization algorithm known for its robustness and high success rate in large parameter systems [41]. Its second-order moment estimation term is particularly useful in solving physics-based optimization problems [42]. In practice, we connect the gradients from `dolfin-adjoint` with the Adam implementation in the machine learning library `PyTorch` [43] using `torch-fenics` [44]. `torch-fenics` enables seamless updating of FE parameters using the Adam optimizer in `PyTorch`.

We previously hypothesized [31] that displacement data should suffice to recover the spatially varying material properties of a target system. Hence, we optimize for the match between target (u_{target}) and optimized (u_{optim}) displacement fields, defined as follows:

$$\mathcal{L}_u = \log_{10} \left(\frac{1}{n} \sqrt{\Sigma_{i=1}^n (u_{target}^i - u_{optim}^i)^2} \right) \tag{3}$$

The formula within the $\log_{10}$ is the root mean squared error (RMSE), with the log term stabilizing the loss function's behavior on displacements of vastly different scales. For instance, in a cantilever beam, the elements at the free end have significantly larger displacements than those near the clamped boundary. This transformation prevents the gradients of large displacements from dominating the gradient descent, leaving regions with relatively smaller displacements unchanged.

2.3. Regularizations

The proposed inverse problem is considered ill-posed, with multiple possible solutions. Regularization helps mitigate ill-posedness, ensuring attention to parameters' search space, not just fitting the model's response to the data. It also helps manage measurement noise, which is crucial since the output of an ill-posed problem can change drastically with minor input variations, potentially destabilizing the optimization process [45,46]. Insufficient data, leading to an infinite number of solutions, is another issue that regularization addresses [47].

To mitigate the ill-posedness and enhance the solution's robustness, we add a regularization term. Common terms in SHM literature include L_1 sparsity-promoting regularization and Tikhonov L_2 regularization [13,48]. We choose L_1 regularization to limit the number of

identified damages. Finding the right balance between the objective and regularization losses through the regularization coefficient α can be challenging.

Through empirical evaluation, we have identified $\alpha = 0.1$ as a suitable compromise. To determine α, we scale the penalty value to match the loss term's order, then test coefficients $\alpha \in \{10^{-4}, 10^1\}$ to find the most appropriate value. Our experiments indicate that $\alpha = 0.1$ effectively penalizes degenerate solutions while optimizing the objective function.

We add the effect of L_1 penalty as the following term to the above loss function:

$$\mathcal{L}_{penalty} = \|\theta_{undamaged} - \theta_{optim}\|_1 \tag{4}$$

$$\mathcal{L}_{final} = \alpha_u + \lambda \mathcal{L}_{penalty} \tag{5}$$

In this context, $\theta_{undamaged}$ represents the homogeneous material state (i.e., no damage), such that $\theta(x) = 1$ everywhere.

We also consider the convergence and the number of iterations required for the inverse solver. Empirically, we have found that allowing the L_1 loss to converge to a near-zero value is crucial for finding the correct solution. Additionally, L_1 loss can serve as a convergence criterion, given that there is no lower bound when optimizing for displacement loss.

In our study, we always use the Adam optimizer with a learning rate of 5×10^{-4} and run the optimizer for 1000 iterations.

2.4. Evaluation metric

Conventionally, an evaluation metric like element-wise norm can be used to compare two solutions, especially when the spatial difference between the results is not a major concern, meaning the proximity or distance of the damaged element from its original location is not crucial. However, in the case of local FEMU, and due to the high resolution of damage localization, it becomes essential to consider adjacency information related to the target and identified damage locations.

Earth Mover's Distance (EMD), also known as the Wasserstein distance, is a statistical measure defined as the minimum cost required to transform one distribution into another. In the context of damage localization, this will be applied to compare the distribution of $\theta(x)$ in space. Unlike the MSE metric we have used in our past work [31], EMD provides a more sophisticated measure that is robust to translations. An illustrative example of the difference between MSE and EMD is two completely disjoint distributions: MSE yields the same value no matter where the samples are, whereas EMD is sensitive to how far apart the samples are.

Our distance metric is defined in the standard way:

$$\mathrm{EMD}(\theta_{GT}, \theta_{optim}) = \frac{\sum_{i=1}^m \sum_{j=1}^n \pi_{i,j} d_{i,j}}{\sum_{i=1}^m \sum_{j=1}^n \pi_{i,j}}, \tag{6}$$

where θ_{GT} is the ground truth fractional elastic modulus, θ_{optim} is the optimized fractional elastic modulus, $D = [d_{i,j}]$ is the distance matrix between elements i in θ_{GT} and j in θ_{optim}, and $\Pi = [\pi_{i,j}]$ denotes the optimal transport plan between these two distributions, subject to minimizing the total cost,

$$cost = \min_\Pi \sum_{i=1}^m \sum_{j=1}^n \pi_{i,j} d_{i,j} \tag{7}$$

In this work, we use the Python Optimal Transport (`POT`) library [49] to solve this minimization problem.

We define the cost matrix D as the Euclidean distance between the centroids of the ground truth and optimized mesh. A convenient feature of this choice is that the EMD has units of length (m) and thus represents how far "mass" (in this case, more like stiffness) must typically be moved to transform the obtained solution into the ground truth. Thus, we report EMD values in m throughout this work. For simplicity, we use the same ground truth and optimized mesh shown in Fig. 2.

Fig. 1. Schematic depiction of the inversion process: The inverse solver F^{-1} computes the optimized fractional elastic modulus for a ground truth (GT) damaged element i, denoted by θ_{optim}^i. Inputs to the solver include the target displacement field u_{target}^i and the initial FE mesh θ_{raw}, which incorporates its boundary and loading conditions. This procedure is performed for each element i to evaluate the algorithm's sensitivity to each damaged element individually throughout the mesh.

3. Results and discussion

3.1. Data generation

To evaluate the usefulness of this approach, we generate a synthetic dataset using FEA. We generate displacement fields using a ground truth model, which is later used as a target for the inverse problem. The process for generating u_{target} is illustrated in Fig. 3 for a simplified mesh. We begin with the raw mesh, θ_{raw}, and for each element in the mesh, we assign a damaged value θ_{dmg} to the ith element while keeping the rest undamaged, θ_{undmg}. In Fig. 3, there are 4 elements, resulting in 4 unique cases θ_{GT}^i, where i denotes the element number and GT indicates that it is the ground truth θ that we will later attempt to recover. Although it is possible to assign multiple damages to θ_{raw}, for demonstration purposes, we focus on single-damage scenarios where only one element is damaged at a time.

For each damage case (four here), we run the physics solver in the forward mode to obtain the displacements at the nodes. This process yields four different u_{target}, which will later be used to recover their corresponding θ_{GT}^i. After calculating the target displacement fields for each u_{target}, we run our inverse solver to attempt to recover their corresponding θ_{GT}^i. This process is depicted in Fig. 1. The inverse solver takes in the initial guess (θ_{raw}) and a target displacement (u_{target}), then produces an optimized solution θ_{optim} through gradient descent.

Next, we calculate the EMD_i between θ_{GT}^i and θ_{optim}^i. First, we compute the Euclidean distance between the two meshes, here assuming the same mesh for θ_{GT} and θ_{optim}, as shown in Fig. 2. We then calculate the EMD_i for each element in the FE mesh θ_{raw}, as illustrated in Fig. 4. For visualization purposes, we display the corresponding EMD value for each element i in the mesh; it is important to note that these represent one inverse solve for each element and are not obtained all at once.

In the following sections, we will consider different scenarios related to generating various displacement fields and recovering the corresponding θ that resulted in u_{target}. Specifically, we define the problem as recovering the spatially varying $\theta(x)$ of an inhomogeneous 2D frame under plane stress. As shown in Fig. 5, the frame is clamped on its left and right ends, and a uniform downward traction load $\omega = -1000$ N/m is applied to the top span. During mesh generation, we ensure symmetrical meshing along the y-axis, halfway along the span ($x = 3$ m).

Our experiments are performed in a high-performance computing environment, typically utilizing 8 CPU cores and up to 32 GB of RAM for each experiment. On average, it takes about 40 s to perform a single inverse solve, and this duration scales with the number of damage cases, which is equivalent to the number of elements in the mesh in most scenarios presented below.

3.2. Single-cell damage localization

In this scenario, we vary the elastic modulus for each element in the mesh, designating one element as damaged ($\theta_{dmg} = 0.5$), while the rest remain undamaged ($\theta_{undmg} = 1$). We then calculate the EMD for every solution the solver finds and assign its value to the corresponding element, as shown in Fig. 5.

To understand how EMD characterizes the solution, we included examples with varying EMD values in Fig. 6. To quantitatively differentiate between successfully localized and mislocalized damage, we establish an EMD threshold of 0.015 m. This threshold is determined by calculating the EMD in a scenario where a single element is damaged, but the solver incorrectly predicts no damage (or vice versa). Specifically, any EMD value below 0.015 m indicates successful damage localization, while values equal to or larger than 0.015 m are classified as mislocalized cases. Note that, here our defined EMD threshold is set to measure the exact damage localization accuracy by defining it as the misclassification of one element. However, a larger EMD may be also acceptable, for example, the second and third row of Fig. 6 show cases where mislocalization happens in the vicinity of the original damage location, and for some use cases, this tolerance is acceptable. Based on the 0.015 m threshold, the solver accuracy is 86.76%, with 118 out of 136 elements. It is important to note that EMD reflects both localization and quantification accuracy, so higher values can mean either locating the wrong damaged element, identifying the wrong modulus magnitude, or both.

As shown in Fig. 5, the EMD in elements near the boundary conditions (such as the lower end of columns and the top span) is relatively low, while the EMD in mid-column elements is higher. This can be physically interpreted as the varying response of different parts of the structure to imposed loads and support conditions; the stability provided by boundary conditions, consistent stress distributions in the beam, and the complex load paths could explain the observed EMD values and the solver's performance in localizing damage. Consequently, there are some blind spots in the columns near the boundary, where the displacement is not sensitive enough to the elastic modulus to accurately detect damaged elements in this region.

3.3. Damage size

Next, we study how the damage size affects localization accuracy. We consider circular damage regions of radius R, such that all cells encompassed within that radius are damaged, rather than a single cell. Note that when R is larger than the smallest characteristic length in the geometry (i.e., the column width), the geometry deviates from a circle since the frame is non-convex. The results in Fig. 7 indicate that as the damage radius increases, the EMD also increases ($r = 0.750$, $p < 0.001$). However, the EMD should be proportional to the size of the damaged region as more cells are included in the calculation ($r = 0.773$, $p < 0.001$). These strong correlations confirm that larger damage regions lead to proportionally higher EMD values, shown on the secondary axis of Fig. 7. The relative error (EMD/Radius ratio) remains fairly consistent for most damage sizes, with means ranging from 0.028 to 0.069. A slight decrease in accuracy is observed only for the largest damage sizes ($R \geq 1.8$ m, relative error ≈ 0.067). This is likely due to the increased complexity of damage spanning multiple structural elements as well as the fewer number of cases tested, mainly because larger damage areas tend to cover all the elements in the frame with fewer cases. This shows no significant change in the quality and accuracy of solutions when moving from a single damaged element to many damaged elements. Overall, these results indicate that the solver maintains a reliable accuracy across different damage sizes, particularly for the well-sampled smaller and medium damage radius scenarios.
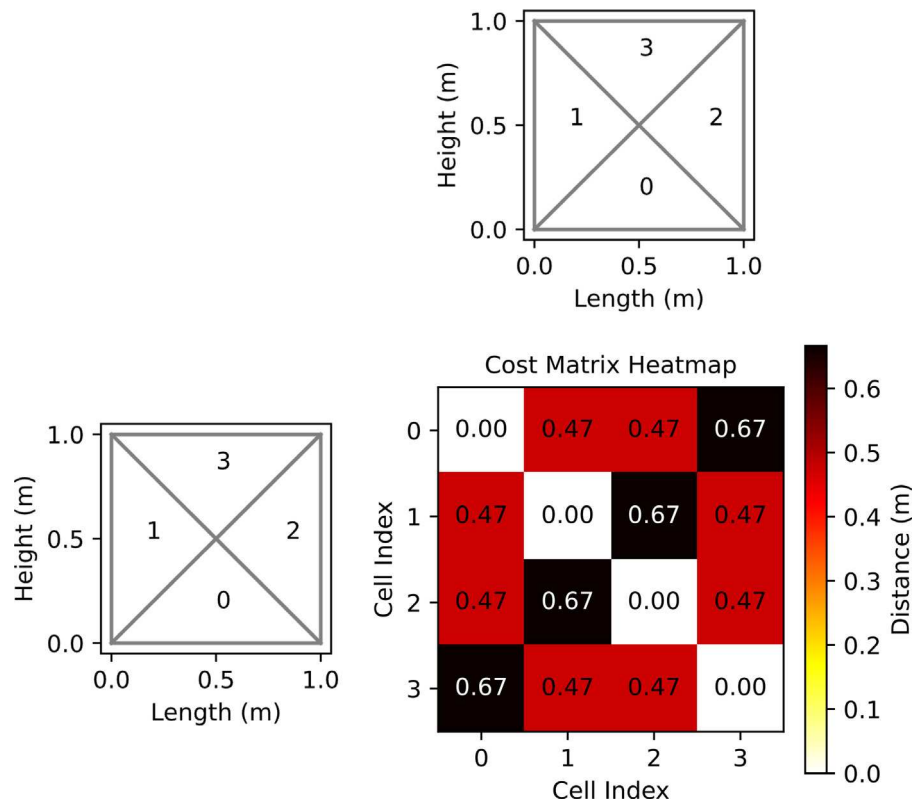
Fig. 2. Illustration of cost matrix calculations used in the evaluation metric. Each element's index is displayed at the element's centroid.
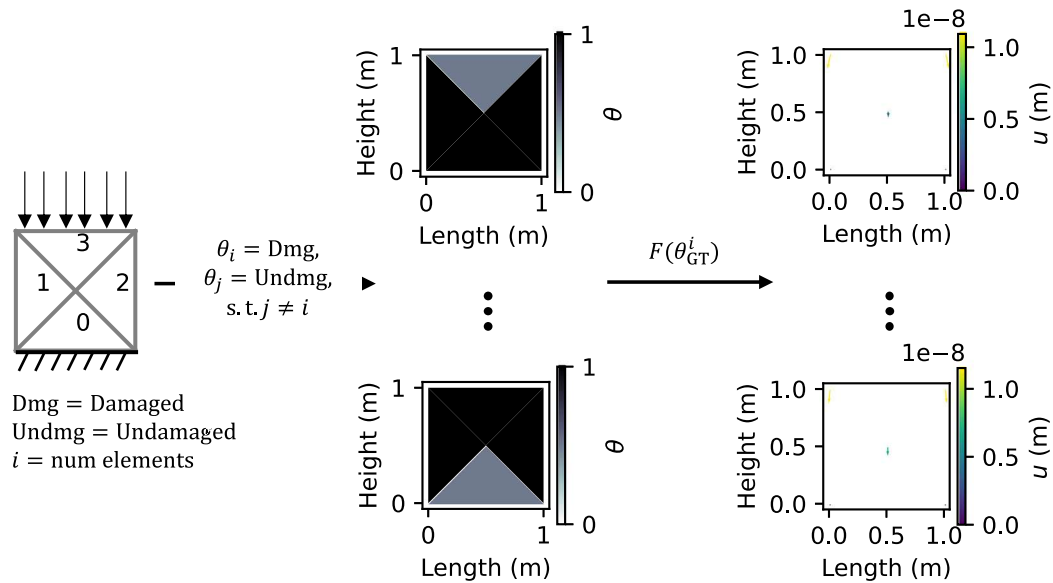

Fig. 3. Process of generating target displacement fields u_{target}: Using the forward linear elastic solver (F) to calculate u_{target} for the damaged ith element, θ_{GT}^i.

3.4. Multi-cell damage localization

In real-world scenarios, a structure may have multiple disconnected damaged elements throughout the system rather than a single damaged element of varying size. Accurately identifying the locations of all key damages is crucial for effective damage assessment. To evaluate the solver's performance in such scenarios, we systematically varied the number of damaged elements (n_d) sequentially: 2, 3, 5, and 10,

testing 544 total cases. For each scenario, an ensemble of n_d random elements is damaged rather than exhaustively enumerating all possible combinations. We report the average EMD for cells that have been damaged across different experiments.

Using our established threshold of 0.015 m for successful localization, the solver achieves success rates of 69.1%, 72.8%, 73.5%, and 86.0% for n_d = 2, 3, 5, and 10 respectively. However, as shown in Fig. 8, increasing the number of damaged elements in the system shifts

Fig. 4. Process of generating EMD map. We assign the EMD to the corresponding element in the FE mesh θ_{raw}. EMD is calculated between the actual damaged case θ_{GT}^i and the recovered case θ_{optim}^i.

the median EMD towards zero, while the number of outliers tends to increase, especially for $n_d = 5$ and 10. This suggests that the solver may not reliably detect more than 5 (disconnected) damaged cells at once. It is worth noting that the penalty term can be adjusted and varied to identify fewer or more damaged elements; while we use a fixed penalty value $\alpha = 0.1$, the choice of varying the penalty term is determined based on one's need and can be used to optimize the trade-off between detection sensitivity and false positives based on specific application requirements.

3.5. Effect of initialization

It is imperative to assess the sensitivity of the gradient-based optimization algorithm to the initial guess used to seed the system. To this end, we start the optimization from different homogeneous states where $\theta_{init} \in \{0.1, 0.2, \ldots, 1.0\}$. We then perform the optimization while keeping the rest of the parameters unchanged. In addition to the homogeneous state, we attempt to start initialization from a noisy state by adding 1σ white Gaussian noise to the half ($\theta_{init} = 0.5$) initialization. However, the solutions show many cells as damaged ($\theta_{optim} < 1$), which is inaccurate. Since that strategy is ineffective, we omit those results.

In Fig. 9, we observe a downward trend in EMD, suggesting that as the initial value increases, the error decreases. We quantify this observation by a Pearson correlation coefficient of ($r = -0.647, p = 0.043$), indicating a significant negative correlation between θ_{init} and EMD. We also do not observe any special preference to $\theta_{init} = 0.5$ where $\theta_{dmg} = \theta_{init}$, suggesting the procedure used throughout our other tests is relatively robust to the initialization strategy.

To benchmark, we also plot the 0.015 m threshold indicated by a horizontal red line in Fig. 9. The line shows that initializing θ outside the range $0.4 \le \theta_{init} \le 0.9$ prevents the solver from converging to useful values. This is likely due to the gradient descent getting stuck in one of the many local minima as it approaches the solution from a poor initial guess. Within the optimal initialization range, mean EMD < threshold, the solver achieves a localization accuracy of 89% with the mean EMD of 0.0092 m.

We also tested the scenario where, for each initial value, we assume its ground truth has no damaged element. As shown in Fig. 10, when the initial value is too small, the optimizer tends to converge to poor solutions that do not reflect the ground truth. As the initial value exceeds 0.4, the mean θ converges to the correct value of 1.0 (per our definition of θ). We are only testing for a homogeneous (undamaged) case; therefore there is only one experiment for each initial value.

We performed further analysis across different ranges of initial value. For low initialization values ($\theta_{init} < 0.4$) the median EMD is significantly higher (0.0263 m), indicating poor performance. For mid-range initializations ($0.4 \le \theta_{init} \le 0.9$), the solver performs best with the lowest median EMD (0.0079 m). High initialization values ($\theta_{init} \ge 0.9$) result in slightly increased median EMD (0.0138 m), which approaches our established threshold of 0.015 m.

From these experiments, we learn that an initial value that is too low, too high, or inhomogeneous can cause the optimizer to converge to incorrect solutions due to the prevalence of local minima in the optimization landscape. Conversely, initializing using a homogeneous approach and within the range of 0.4 to 0.9 provides a good balance, ensuring reliable convergence to accurate solutions. Our strategy using a mid-range homogeneous initialization $\theta_{init} = 0.5$ strikes this balance effectively, demonstrating robustness in the optimization process.

3.6. Mesh resolution

We next evaluate the solver's sensitivity to mesh resolution by analyzing the relationship between element size and localization accuracy. In other words, we probe the smallest damage we can detect using this displacement-based inverse solver. For this matter, we use a fixed-fixed beam geometry, which allows for better control over the mesh resolution as shown in Fig. 12 (top). We can now vary the number of elements in the subsequent experiment.

In Fig. 13, we present the effect of increasing mesh resolution. We observe a corresponding decrease in the median EMD distribution for each case. The resulting median EMD for each case were 1.05×10^{-2} m, 2.65×10^{-3} m, 5.43×10^{-4} m, 1.50×10^{-4} m for 64, 256, 1024, and 4096 element beams, respectively. This represents a 3.97, 4.87, and 3.62-fold decrease in median EMD values as the number of elements increased.

To explain this trend, we plot the average size of the elements in each mesh, which is equivalent to the damage size in the single-element damage detection scenario. The data indicate that the median EMD is directly proportional to the damage size ($r = 0.9999$, $p < 0.001$). However, we also observed that the coefficient of variation (CV) increases as the mesh gets finer. Specifically, the CV values are 20.3% for 64 and 256 element beams, 41.7% for 1024 element beam, and 34.3% for 4096 element beam. This increased variability suggests localization becomes more challenging for finer mesh.

Localization accuracy was highest, with the lowest CV, for 64 and 256 element beams. The corresponding single-element damage area is approximately 6.25×10^{-2} m^2 (i.e., 25 cm × 25 cm) and 1.56×10^{-2} m^2 (roughly 12 cm × 12 cm), respectively. In contrast, Fig. 11 shows that localization becomes increasingly challenging with finer meshes, such as 1024 element beam with an element size of 3.90×10^{-3} m^2 (roughly 5 cm × 5 cm) and 4096 element beam with an element size of 9.76×10^{-4} m^2 (approximately 3 cm × 3 cm).

The inverse solver shows that detection is possible with finer mesh and the median EMD decreases proportional to the element size while increasing the CV. This reveals that there is a fundamental trade-off: while finer meshes enable the localization of smaller damage features, they also introduce greater variability in localization accuracy.

3.7. Practical considerations

While most of the aforementioned experiments were conducted under the assumption of perfect knowledge of the system and measurement, the following scenarios will probe the effectiveness of the framework in the presence of uncertainties that arise in real-world applications (e.g., modeled after DIC). This includes uncertainties in measurements or mesh geometry. The aim of this section is to further assess the robustness of the proposed method in practice.

3.7.1. Measurement noise

Noise is an inherent part of any real-world experiment. We add various degrees of Gaussian white noise to the displacement fields, allowing for the consideration of noisy measurements in the field. Since we use both *x*- and *y*-displacement components, noise is added to each component separately, according to its distribution. The following noise levels are added to the system: {0.5%, 1%, 2%, 4%, 8%, 16%}. As suggested in a recent study, the measured displacement by the 3DIC camera had an average of about 3% noise in its values [11].
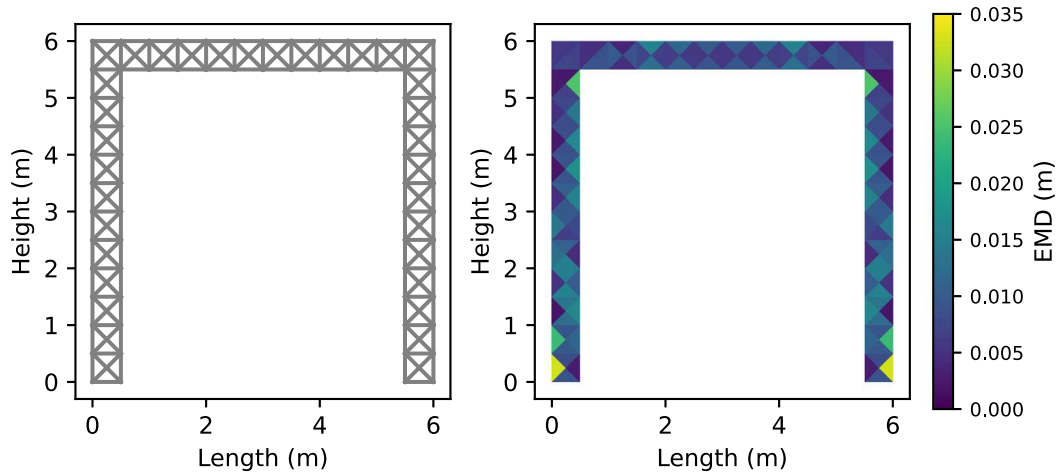
Fig. 5. (Left) Visualization of θ_{raw} with 136 elements, subject to a downward distributed load ω. (Right) EMD map of the frame.

We find that as the noise level increases, the median EMD approaches the 0.015 m threshold value (indicated again by a red line in Fig. 14). With localization accuracy declining from 83.8% at 0% noise to 71.32% at 2% noise to 59.6% at 16% noise. Particularly, the mean EMD value in low-noise conditions ($\leq$ 2%) 0.0016 m (std = 0.0040 m) degrades significantly in high-noise conditions ($\geq$ 2%) 0.0142 m (std = 0.0020 m). Dizaji et al. [11] report a similar drop in range 5%–10% in the localization and detection accuracy when transitioning from simulation to experiment.

Detailed analysis shows that the localization error often starts in the columns and gradually expands to the beam as noise increases, with column EMD values ranging from 0.0105 m to 0.0150 m compared to beam EMD values of 0.0078 m to 0.0132 m. This differential sensitivity can be potentially explained by the distributed loading on the beam imposing noticeable deformation compared to the columns, making it robust to small measurement variations [50]. Therefore, smaller displacements in columns make it challenging to diagnose them as noise increases. These findings suggest that measurements with noise levels beyond 2% should be carefully studied, and multiple loading conditions might help to achieve a balanced solution that is robust to noise, particularly for column damage identification.

3.7.2. Reconstruction from limited measurements

Here we examine how the quality of the solution is impacted by reducing the measurements to only a portion of the full field, emulating the effect of surface-level measurement (e.g., DIC) as opposed to subsurface nondestructive evaluation techniques. Fig. 15 shows the variation in displacement measurement density with changes in penetration depth (ϵ). This depth-dependent density of displacement measurements enables a comparison between surface and full-field measurements. Note that for this experiment, we need to work with thicker elements (beam and columns) compared to the frame in preceding results, so that increasing the penetration depth adds a significant number of observations for the solver.

We start the experiment by using the surface displacement $\epsilon = 0.1$ m (116 vertices) and test with 0.3 m (229 vertices) and 0.5 m (284 vertices). Each vertex has an x- and y-component that will be used in the inversion process. Also, note that the number of observed points in $\epsilon = 0.5$ m is equivalent to the full-field case.

Fig. 16 shows that median EMD value improves modestly by increasing penetration depth: 2.08×10^{-3} m at $\epsilon = 0.1$ m, 1.97×10^{-3} m at $\epsilon = 0.3$ m, and 1.93×10^{-3} m at $\epsilon = 0.5$ m. The localization accuracy, defined as the percentage of elements below EMD < 0.0039 m remains consistent: 82.9% for both surface- and intermediate-level depths and a slight increase to 84.2% using the full-depth information. The 0.0039 m threshold was calculated for this frame and defined, as before, as the

EMD value of a single-damaged frame to an undamaged frame.

In this experiment, surface measurements capture most of the information in the full field, suggesting that increasing the measurement depth does not systematically affect the EMD value (r = −0.001, p = 0.725). While there is an increase in the number of correctly localized elements as the penetration depth ϵ increases, however, the relatively similar median and EMD scores suggest that there is not much additional information that can be gleaned from the full-field displacement.

Our results support the use of surface deformations as a proxy to full-field measurement in this scenario, achieving 82.9% accuracy with only 40.8% of the measurement points. This observation is on par with the literature where the authors report robust localization and detection of subsurface damage using DIC surface measurement [11]. However, the results should be evaluated for specific problem requirements to be used in non-destructive schemes.

3.7.3. Problem geometry and size

So far, we have studied a single-story, single-span frame for damage localization, as shown in Fig. 5 (left). To demonstrate that our algorithm is robust against changing mesh geometry and size, we now investigate two additional scenarios: a fixed-fixed beam and a multi-span multi-story frame, as shown in Fig. 17. The columns are fixed at their lowest part (y = 0 m).

Due to the increased number of elements, this two-story and three-span frame is a more difficult inverse problem. To control the effect of variation in detection resolution (average element size), we included results for a fixed-fixed beam with 440 elements, a single-story frame (SSF) with 450 elements, and a multi-story frame (MSF) with 540 elements.

We assume a distributed downward vertical loading on the top and second floor of the MSF, resulting in the shown EMD map in Fig. 17 (bottom), where the elements on the first and second floors have relatively low EMD.

We notice that elements with larger EMD are mostly located near the edges and corners, perhaps due to the complex loading and stress patterns in these areas. These locations, where beams and columns intersect, experience complex stress states because of the interaction of multiple force transfer mechanisms such as bending moments, shear transfer, and geometric discontinuity. This complexity poses a challenge for the inverse solver to accurately disentangle and back-calculate the sensitivity of multiple interacting elements, ultimately leading to reduced localization accuracy in these regions.

Additionally, we examine a beam with two fixed ends at x = 0 m and 4 m, shown in Fig. 17 (top). The beam is under 1 kN/m downward point loading. Compared to both the single-story frame and MSF, Fig. 18

Fig. 6. Comparative examples of small, medium, and large EMD values. Solutions are zoomed-in to aid with visualizing the damage location. Ground Truth, θ_{GT}, shown on the left and optimized, θ_{optim}, on the right.

shows consistently lower EMDs (median = 0.00076 m) for the beam. This is compared to both the SSF (median = 0.0019 m) and MSF (median = 0.0030 m). The higher localization accuracy of the beam case can be attributed to its simpler geometry and straightforward load path. However, even in this simpler case, we find that damages closer to the beam's boundary are less likely to be accurately localized, i.e., have a higher EMD. This observation is consistent with what we observed with the SSF and MSF, where damages in low displacement regions are less likely to be identified.

We observe a significant difference in EMD distributions across the three structures. The results show that increasing the geometrical complexity influences the localization accuracy, with the largest disparity observed between the beam and MSF.

3.7.4. Measurement and mesh error tolerance

Another limiting assumption of this work is that we generate the target displacement field and run the inverse solver using the same mesh. This assumption is problematic in practice, as measurements of real-world objects are discrete. In such cases, the Finite Element (FE) model

Fig. 7. EMD distributions for varying damage radii, incrementally increasing by enlarging the radius R and damaging adjacent elements around the arbitrarily chosen initial damaged element's centroid.
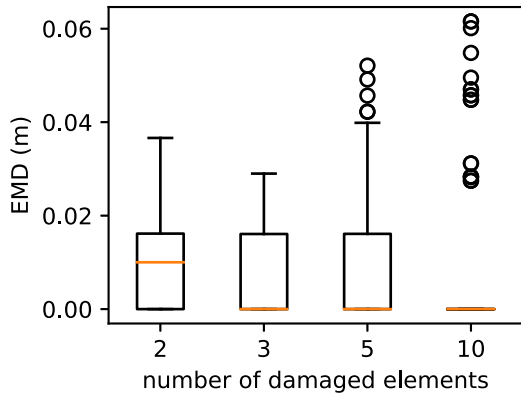

**Fig. 8.** EMD distributions for cases with multiple simultaneous damages across the system.

Fig. 9. EMD distributions as a function varying initial values, starting from 0.1 and incrementally increasing by 0.1 steps to reach $\theta_{\max} = 1.0$.

often has more elements/nodes than the measurements. Conversely, DIC systems provide high-resolution measurements of the displacement field, where the FE model may have fewer elements than the number of measurements. We emulate the effects of resolution mismatch between the measured data and the model in two different experiments here to investigate the implications in real-world scenarios.

As shown in Fig. 19 (left), we first investigate the conventional approach, where the measured data is discrete, thus providing fewer measurements to reconstruct the solution. For this, we use 136 elements (104 nodes) to generate the target displacement and recover the
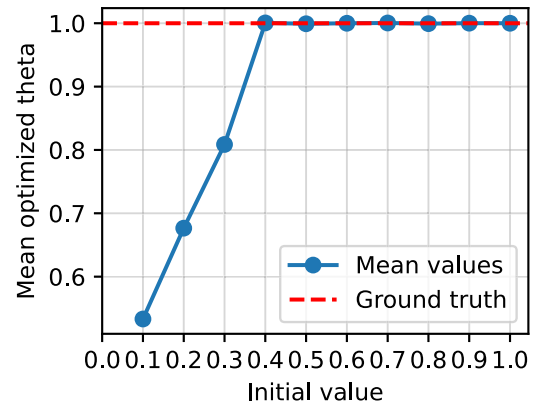

Fig. 10. Mean values of optimized structure found with different initial values when there is no damage in the frame.

solution with 241 elements (181 nodes). Note that the total number of measurements of the target displacement field is 104×2 because we measure the x- and y-displacement for each node. For the optimization process, we update the FE mesh with 241 elements, meaning we use 208 measurements to recover 241 unknowns. The discrepancy between the existing elements and nodes arises because we use a triangular mesh, leading to some overlapping and non-overlapping nodes.

Fig. 19 (left) shows the result for the conventional approach, suggesting that we can recover damage existing in all regions of the original frame (used for obtaining measurements) with high accuracy. We use the ground truth mesh to display the solutions because the original damaged elements were assigned to it.

Secondly, we investigate a case more relevant to our argument in this work based on the idea of DIC data collection. As shown in Fig. 20 (right), in the case where the measurement resolution is higher than the FE mesh used for recovering θ, we still achieve a comparable EMD to the traditional case. We use 362 measurements to recover 136 unknowns. Additionally, by comparing the EMDs in Fig. 20, we observe that the median EMDs are comparable for both scenarios, at 0.1334 ± 0.0056 m for the upsampling (fewer measurement) case and 0.1344 ± 0.0030 m for the downsampling (more measurement) case. However, a Kolmogorov–Smirnov test shows that the EMD distributions are statistically different ($D = 0.248$, $p < 0.001$), indicating that while median EMD values are comparable, the resulting distributions differ from each other, such as differences in the spread and the tail, as seen in Fig. 20.

This section focused on whether a mismatch between the GT and optimized mesh results in significantly lower damage localization accuracy. According to Fig. 20, there is no noticeable change in the median EMDs, indicating that the solver is robust to a small mismatch between the number of measurements and unknowns in the FE model at the tested scale. The results are promising for the tested problem. However, it should be noted that as long as the number of linearly independent measurements is larger than the number of unknowns, we will arrive at a unique solution. If this condition is violated, we may have a potentially infinite number of solutions [5].

3.7.5. Measurement mismatch

We finally examine the implications of using displacement measurements measured and evaluated at imprecise nodal points of the FE model. This assumption introduces a notable limitation, particularly when the actual measurement locations do not perfectly align with the FE model's nodal coordinates. To assess the impact of this misalignment, we introduce Gaussian white noise with a standard deviation of 1σ to the x- and y-components of nodal coordinates. The standard deviation values, σ, are varied across a range: 0.5%, 1%, 2%, 4%, and 8%.
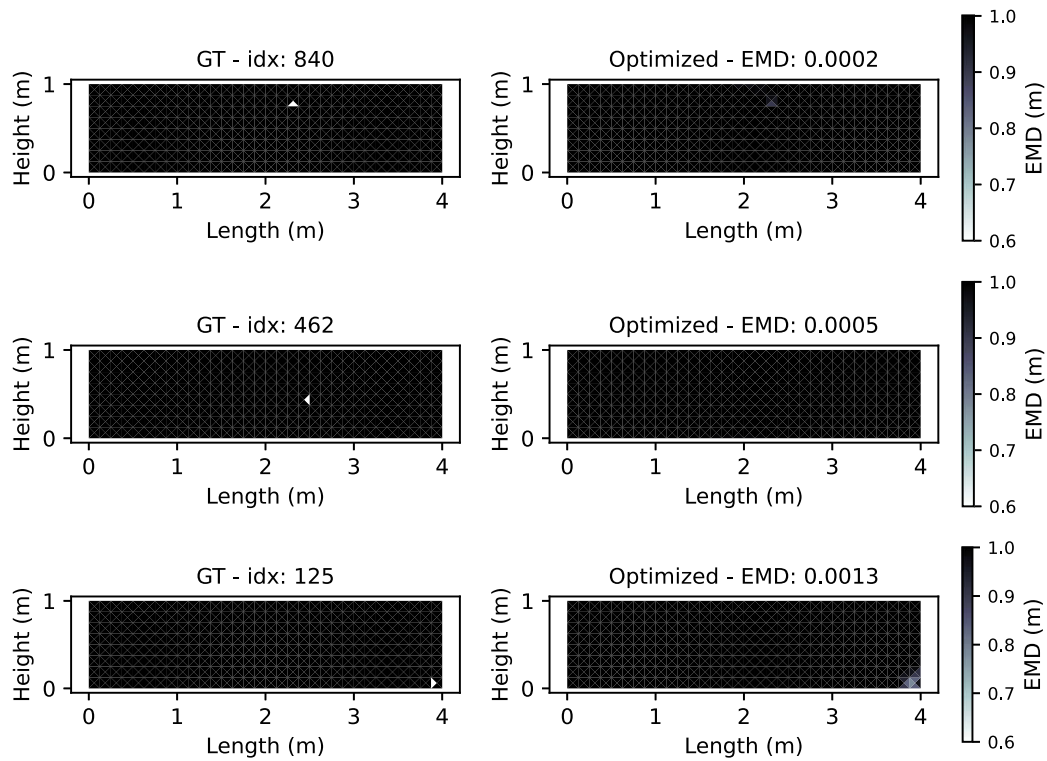
Fig. 11. Examples of small, medium, and large EMD values for the fixed-fixed beam consisting of 1024 elements. Solutions are zoomed in to aid with visualizing the damage location. Ground Truth, θ_{GT}, shown on the left and optimized, θ_{optim}, on the right.
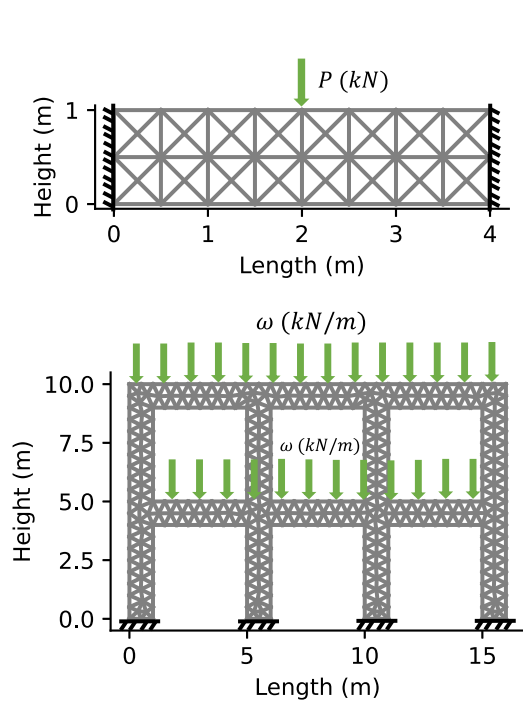

Fig. 12. Illustration of (top) a fixed-fixed beam consisting of 64 elements, subjected to a downward point load, P. (bottom) a multi-story frame (MSF) consisting of 540 elements subject to downward distributed load ω on the second and first floors.

Fig. 13. EMD distributions for the fixed-fixed beam in Fig. 12 as the number of elements in the model increases.

Fig. 14. EMD distributions in response to increasing input measurement noise, incremented in multiples of 2 (e.g., 1, 2, 4, 8, 16).
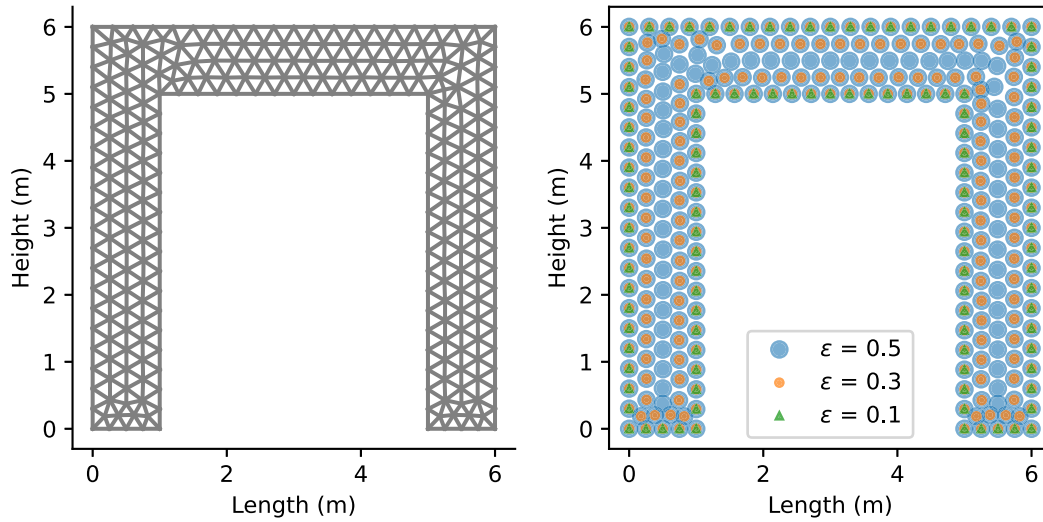
Fig. 15. Illustration of the effect of penetration depth on selected vertices. The sampled vertices are color-coded to reflect various levels of penetration depth. As the penetration becomes deeper, the number of highlighted sampling points increases, indicating a greater extent of measured displacement.


Fig. 16. EMD distributions as a function of varying penetration depth ϵ. As ϵ increases, the number of measurements approaches the observable full-field data.

Sampling at the perturbed coordinates can be challenging as some coordinates may fall outside the predefined geometry of the FE mesh. To address this, we recalibrate components that deviate beyond the geometry, leading to new noise percentages of $0.29\%, 0.62\%, 1.13\%, 2.33\%, 3.85\%, 6.51\%$. Fig. 21 shows that as the mismatch noise increases, the median shifts towards the baseline red line, indicating that fewer elements are localized accurately. Quantitatively, at 0.29% noise, the localization accuracy is 86.03%. As the noise level increases to 2.33% noise, the localization accuracy drops to 72.06% and remains fairly consistent thereafter.

Results imply that for noise levels above 2%, the solver tends to converge to an undamaged solution. Hence, as the coordinates mismatch increases, the results become less reliable, and some may have already diverged.

4. Conclusions

In this study, we demonstrated the use of a differentiable-physics-based inverse solver to localize element-level damage in a single-story frame with 83.3% accuracy. We found that our approach could reliably identify and localize damage at an element-level accuracy with a 72%

accuracy considering measurement and mesh alignment noise. We use the EMD evaluation metric that helped us identify the locations where the framework generates reliable results and identify locations where the solutions may be less certain. Additionally, compared to previous studies that typically modeled frames with few elements [4,50,51], our method is able to accommodate structures with a large number of parameters, here we show frames with up to 540 elements, and beam with up to 4096 elements. This represents a significant advancement in terms of model complexity and resolution.

We evaluated the robustness of the inverse solver under a variety of conditions to emulate real-world use. These tests include different damage configurations, initialization, noise levels in measurements, changes in system geometry, and variations in model resolution. Despite the inherent challenges associated with these conditions, the solver was robust to most of these perturbations. However, it is important to note that real-world noise levels, especially in the 3%–5% range, posed significant challenges. Furthermore, it is essential to note that our methodology and system have primarily been tested with synthetic data. The generated data, while instrumental in providing controlled conditions for initial testing, may not capture important effects present in real-world SHM applications.

While our results are promising, there are several limitations to our approach that need to be acknowledged. Firstly, the study is restricted to damaged elements, rather than discontinuities that might be associated with fracture. Secondly, the linear elasticity assumption inherent in our methodology might be an oversimplification for some structures.

Third, we relied on a local optimizer (gradient descent) in our framework. Although we implemented several techniques to tackle the ill-posedness of the problem, the nature of local optimization means that there is a substantial likelihood of arriving at solutions that diverge from the ground truth. Moreover, the assumption of exhaustive knowledge of loading and boundary conditions can be a potential limitation in real-world scenarios.

Considering these points, we argue that differentiable physics is a promising new tool to augment conventional finite element model updating techniques. It enables fast and scalable inversion of structural systems at the element-level, which pairs well with high-resolution, precise measurements provided by modern full-field sensing techniques.
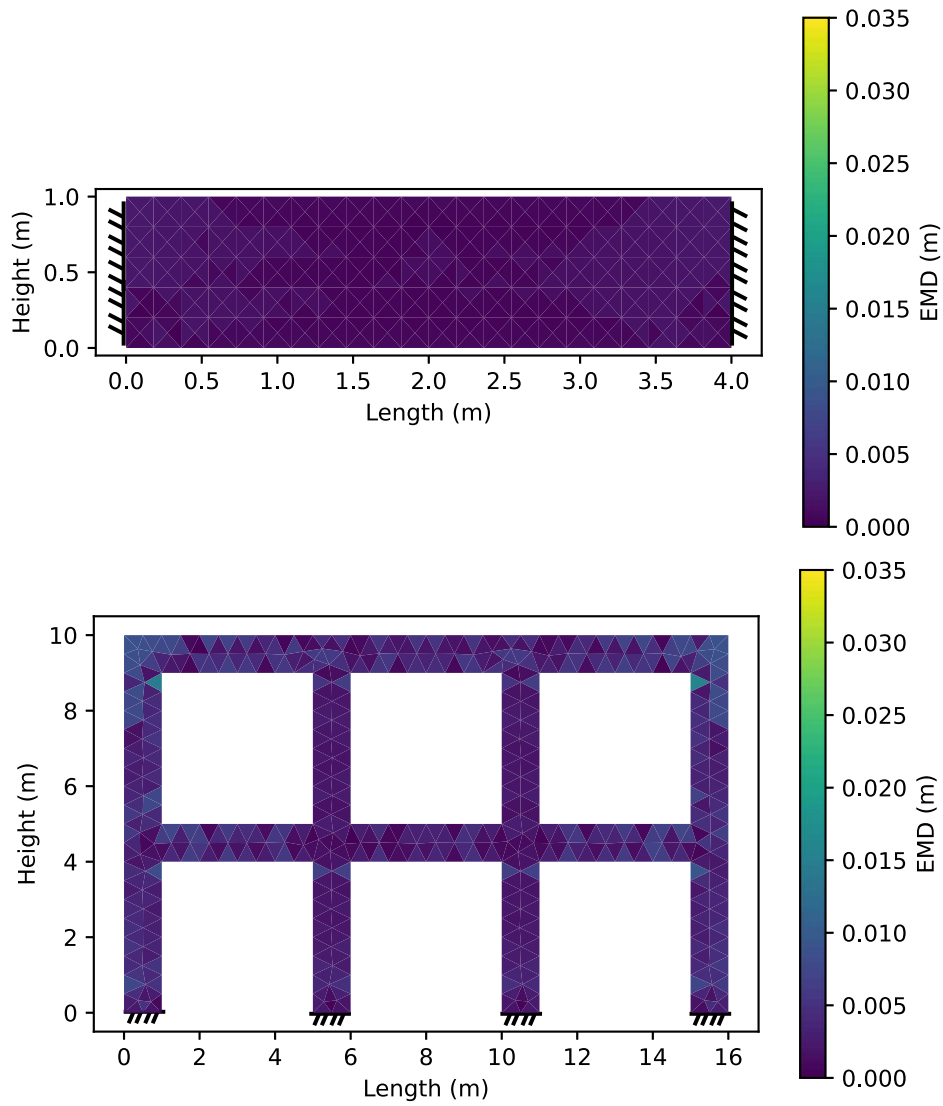
Fig. 17. EMD maps for (top) a fixed-fixed beam subject to a point load and (bottom) the multi-story frame (MSF) subjected to traction loading on the first and second floors.


Fig. 18. Comparison between the EMD distributions of different structural geometries.

Gradient descent makes the solver highly scalable, with the resolution demonstrated here likely being sufficient to provide an initial assessment of a structure's condition. Additionally, our results suggest that the differentiable physics depends only on having a representative, rather than exact, representation of the system.

In the future, we hope to further explore the limitations of the approach and implement it in a controlled laboratory environment. This could involve constructing scaled-down replicas of structural elements like cantilever beams, fixed-fixed beams, and a down-sized single-story single-span frame. Future studies could consider testing for various static loading scenarios where the effects of induced loads on the structure with different directions, magnitudes, and durations can be studied. Additionally, the laboratory setting will allow for precise control over lighting conditions, including light intensity, orientation, and the presence of shadows, which can influence the accuracy of DIC measurements. Furthermore, it might be of interest to investigate the effectiveness of combining dynamic vibration data, such as mode shapes, with static displacement to increase the method's robustness and accuracy. These laboratory-scale experiments play a crucial role in validating the proposed method and understanding its strengths and limitations in practice.

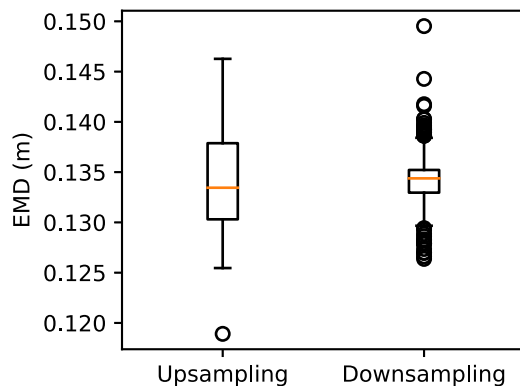**Fig. 19.** Zoomed-in EMD map showing (left) upsampling and (right) downsampling scenarios.

Fig. 20. Comparison of EMD distributions between upsampled (discrete) and downsampled (continuous) measurement scenarios.


Fig. 21. EMD distribution for scenarios with varying degrees of misalignment between the measurement nodes and the constructed FE model.

CRediT authorship contribution statement

Borna Rahnamay Farnod: Writing – review & editing, Writing – original draft, Visualization, Software, Data curation, Conceptualization. **Wesley F. Reinhart:** Writing – review & editing, Validation, Supervision, Resources, Methodology, Funding acquisition, Conceptualization. **Rebecca K. Napolitano:** Writing – review & editing, Validation, Supervision, Funding acquisition, Formal analysis, Conceptualization.

Declaration of competing interest

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

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the author(s) used the GPT4 LLM in order to improve the grammar of the author(s) original writing. After using this tool/service, the author(s) reviewed and edited the content as needed and take(s) full responsibility for the content of the publication.

Acknowledgments

Data availability

Data will be made available on request.

References

[1] Bridges report card for america infrastructure. 2023, https://infrastructurereportcard.org/wp-content/uploads/2020/12/Bridges-2021.pdf. [Accessed 25 December 2023].

[2] Warner James E, Hochhalter Jacob D, Leser William P, Leser Patrick E, Newman John A. A computationally-efficient inverse approach to probabilistic strain-based damage diagnosis. In: Annual conference of the PHM society, vol. 8, 2016.

[3] Cai Jing, Wu ZS. Statistical damage localization and assessment of frame structures from static test data. J Struct Eng 2001;47:9–18.

[4] Amiri HR, Esfandiari MJ, Dehghan Manshadi SH, Urgessa G. A novel sensitivity-based method for damage detection of a structural element. Asian J Civ Eng 2020;21:1079–93.

[5] Sanayei Masoud, Nelson Richard B. Identification of structural element stiffnesses from incomplete static test data. SAE Trans 1986;1237–48.

[6] Farrar Charles R, Worden Keith. Structural health monitoring: a machine learning perspective. John Wiley & Sons; 2012.

[7] Sanayei Masoud, Saletnik Michael J. Parameter estimation of structures from static strain measurements. I: Formulation. J Struct Eng 1996;122(5):555–62.

[8] Banan Mo R, Banan Ma R, Hjelmstad KD. Parameter estimation of structures from static response. I. Computational aspects. J Struct Eng 1994;120(11):3243–58.

[9] Hajela P, Soeiro FJ. Structural damage detection based on static and modal analysis. AIAA J 1990;28(6):1110–5.

[10] Sanayei Masoud, Onipede Oladipo. Damage assessment of structures using static test data. AIAA J 1991;29(7):1174–9.

[11] Dizaji MS, Alipour M, Harris DK. Subsurface damage detection and structural health monitoring using digital image correlation and topology optimization. Eng Struct 2021;230:111712.

[12] Zaletelj Klemen, Slavič Janko, Boltežar Miha. Full-field DIC-based model updating for localized parameter identification. Mech Syst Signal Process 2022;164:108287.

[13] Li Li. Damage and material identification using inverse analysis. Electron Theses Diss 2012;4825.

[14] Zanarini Alessandro. Full field optical measurements in experimental modal analysis and model updating. J Sound Vib 2019;442:817–42.

[15] Reu Phillip L, Sweatt William, Miller Timothy, Fleming Darryn. Camera system resolution and its influence on digital image correlation. Exp Mech 2015;55:9–25.

[16] Seidl D Thomas, Granzow Brian N. Calibration of elastoplastic constitutive model parameters from full-field data with automatic differentiation-based sensitivities. Internat J Numer Methods Engrg 2022;123(1):69–100.

[17] Claire Damien, Hild François, Roux Stéphane. A finite element formulation to identify damage fields: the equilibrium gap method. Int J Numer Methods Eng 2004;61(2):189–208.

[18] Hild François, Roux Stéphane. Digital image correlation: from displacement measurement to identification of elastic properties–a review. Strain 2006;42(2):69–80.

[19] Al-Ruzouq Rami, Dabous Saleh Abu, Junaid M Talha, Hosny Fatma. Nondestructive deformation measurements and crack assessment of concrete structure using close-range photogrammetry. Results Eng 2023;18:101058.

[20] Teo T-A. 3D deformation measurement of concrete wall using close-range photogrammetry. Int Arch Photogramm Remote Sens Spatial Inf Sci 2020;43:1175–9.

[21] Liu Yanhua, Zeng Lei, Xiang Sheng, Mo Jinxu, Zhang Jicheng, Chen Juan, et al. Compressive performance evaluation of concrete confined by stirrups at elevated temperature using DIC technology. Constr Build Mater 2020;260:119883.

[22] Hu Jun, Liu Ensheng, Yu Jiayu. Application of structural deformation monitoring based on close-range photogrammetry technology. Adv Civ Eng 2021;2021:1–11.

[23] Ma Qiaoyu, Solís Mario, Galvín Pedro. Wavelet analysis of static deflections for multiple damage identification in beams. Mech Syst Signal Process 2021;147:107103.

[24] Morlier Joseph, Salom Pierre, Bos Frédéric. New image processing tools for structural dynamic monitoring. Key Eng Mater 2007;347:239–44.

[25] Reagan Daniel, Sabato Alessandro, Niezrecki Christopher. Feasibility of using digital image correlation for unmanned aerial vehicle structural health monitoring of bridges. Struct Health Monit 2018;17(5):1056–72.

[26] Wu Rong, Zhang Dongsheng, Yu Qifeng, Jiang Yuxi, Arola Dwayne. Health monitoring of wind turbine blades in operation using three-dimensional digital image correlation. Mech Syst Signal Process 2019;130:470–83.

[27] Molina-Viedma Ángel Jesús, Pieczonka Lukasz, Mendrok Krzysztof, López-Alba Elías, Díaz Francisco A. Damage identification in frame structures using high-speed digital image correlation and local modal filtration. Struct Control Health Monit 2020;27(9):e2586.

[28] Rocchetta Roberto, Broggi Matteo, Huchet Quentin, Patelli Edoardo. On-line Bayesian model updating for structural health monitoring. Mech Syst Signal Process 2018;103:174–95.

[29] Shafiei Dizaji Mehrdad, Alipour Mohamad, Harris Devin K. Leveraging vision for structural identification: a digital image correlation based approach. In: International digital imaging correlation society: proceedings of the first annual conference, 2016. Springer; 2017, p. 121–4.

[30] Shafiei Dizaji M, Alipour M, Harris DK. Leveraging full-field measurement from 3D digital image correlation for structural identification. Exp Mech 2018;58:1049–66.

[31] Farnod Borna Rahnamay, Reinhart Wesley F, Napolitano Rebecca K. Displacement-based structural identification using differentiable physics. Eng Struct 2023;293:116632.

[32] Payawal John Mark Go, Kim Dong-Keon. Image-based structural health monitoring: A systematic review. Appl Sci 2023;13(2):968.

[33] Barros Francisco, Aguiar Susana, Sousa Pedro J, Cachaço António, Tavares Paulo J, Moreira Pedro MGP, et al. Displacement monitoring of a pedestrian bridge using 3D digital image correlation. Procedia Struct Integr 2022;37:880–7.

[34] Eslami SM, Abdollahi F, Shahmiri J, Tavakkoli SM. Structural damage detection by using topology optimization for plane stress problems. Int J Optim Civ Eng 2019;9(1):159–76.

[35] Griewank Andreas, Walther Andrea. Evaluating derivatives: principles and techniques of algorithmic differentiation. SIAM; 2008.

[36] Prasad Rachit, Kim Hyunsoon, Choi Seongim. Adjoint based finite element model updating and validation using time domain based FSI analysis. In: 58th AIAA/aSCE/AHS/aSC structures, structural dynamics, and materials conference. 2017, p. 1125.

[37] Mitusch Sebastian K, Funke Simon W, Dokken Jørgen S. Dolfin-adjoint 2018.1: automated adjoints for ?FEniCS and firedrake. J Open Source Softw 2019;4(38):1292.

[38] Alnæs Martin, Blechta Jan, Hake Johan, Johansson August, Kehlet Benjamin, Logg Anders, et al. The FEniCS project version 1.5. Arch Numer Softw 2015;3(100).

[39] Farnod Borna R, Reinhart Wesley F, Napolitano Rebecca N. Towards damage detection using physics-based modeling and data-driven optimization. In: Sensors and smart structures technologies for civil, mechanical, and aerospace systems 2022, vol. 12046, SPIE; 2022, p. 348–54.

[40] Kingma Diederik P, Ba Jimmy. Adam: A method for stochastic optimization. 2014, arXiv preprint arXiv:1412.6980.

[41] Brown Tom, Mann Benjamin, Ryder Nick, Subbiah Melanie, Kaplan Jared D, Dhariwal Prafulla, et al. Language models are few-shot learners. In: Advances in neural information processing systems, vol. 33, 2020, p. 1877–901.

[42] Thuerey Nils, Holl Philipp, Mueller Maximilian, Schnell Patrick, Trost Felix, Um Kiwon. Physics-based deep learning. WWW; 2021.

[43] Paszke Adam, Gross Sam, Massa Francisco, Lerer Adam, Bradbury James, Chanan Gregory, et al. Pytorch: An imperative style, high-performance deep learning library. In: Advances in neural information processing systems, vol. 32, 2019.

[44] Barkman Patrik. Torch-FEniCS. 2019, https://github.com/barkm/torch-fenics. Accessed 25 December 2023,

[45] Benning Martin, Burger Martin. Modern regularization methods for inverse problems. Acta Numer 2018;27:1–111.

[46] Gallet Adrien, Rigby Samuel, Tallman TN, Kong Xiangxiong, Hajirasouliha Iman, Liew Andrew, et al. Structural engineering from an inverse problems perspective. Proc R Soc A 2022;478(2257):20210526.

[47] Ereiz Suzana, Duvnjak Ivan, Jiménez-Alonso Javier Fernando. Review of finite element model updating methods for structural applications. In: Structures, vol. 41, Elsevier; 2022, p. 684–723.

[48] Fan Xingyu, Li Jun, Hao Hong, Ma Shenglan. Identification of minor structural damage based on electromechanical impedance sensitivity and sparse regularization. J Aerosp Eng 2018;31(5):04018061.

[49] Flamary Rémi, Courty Nicolas, Gramfort Alexandre, Alaya Mokhtar Z, Boisbunon Aurélie, Chambon Stanislas, et al. POT: Python optimal transport. J Mach Learn Res 2021;22(78):1–8.

[50] Bakhtiari-Nejad F, Rahai A, Esfandiari A. A structural damage detection method using static noisy data. Eng Struct 2005;27(12):1784–93.

[51] Rezaiee-Pajand Mohaamad, Kazemiyan Mohammad Sadegh, Aftabi. S A. Static damage identification of 3D and 2D frames. Mech Based Des Struct Mach 2014;42(1):70–96.