

Efficient Structural Reconnaissance Surveying for Regional Postseismic Damage Inference with Optimal Inspection Scheduling

Mohamadreza Sheibani¹; Yinhu Wang²; Ge Ou³; and Nikola Marković⁴

Abstract: Accurately estimating the extent of damage after an earthquake requires labor-intensive reconnaissance surveys, which may take months to cover the entire building inventory in an impacted region. This paper provides a data-driven framework to guide a survey team efficiently through a reconnaissance mission and estimate regionwide damage by inspecting only a fraction of buildings. First, it is shown that by considering a relatively small set of representative buildings in the training data, the necessity of inspecting the entire building inventory is diminished, and accurate estimation of the regional damage is made possible within 2 weeks after the earthquake. Second, to develop a cost-effective solution, the problem of prioritizing buildings and designing efficient inspection routes is formulated as an orienteering problem. The results of the sparse field observations obtained by the end of each inspection day are used to retrain a Gaussian process regression model, which is applied to estimate damage for the uninspected buildings. A regional earthquake simulation testbed was used to validate and evaluate the performance of the proposed method. DOI: [10.1061/\(ASCE\)EM.1943-7889.0002069](https://doi.org/10.1061/(ASCE)EM.1943-7889.0002069). © 2021 American Society of Civil Engineers.

Introduction

The efforts that go into managing essential tasks in the aftermath of a severe earthquake have a significant impact on the pace of recovery. Accurate estimation of building and infrastructure damage improves the efficiency of these efforts. For example, having a rapid and reliable postdisaster needs assessment (PDNA) assists organizations in requesting the appropriate amounts of financial aid, which helps prevent any recovery delays (Loos et al. 2020). Furthermore, identifying the damaged buildings in a timely manner provides valuable information that can be used to determine the reoccupation risks and required recovery tasks (Ural et al. 2011; Lenjani et al. 2020b).

Creating a geographical map featuring the damage intensities of individual buildings has been the emerging focus of various recent studies. In current practice, the extent of building damage is approximated with several methods, including vulnerability assessment, unmanned aerial vehicle (UAV) images, and crowdsourcing from social media (Wilkinson et al. 2018; Bland and Frost 2013). Vulnerability functions commonly are used to provide a probabilistic approach for damage and loss estimation at a regional level (Xin et al.

2018; Steelman and Hajjar 2009). However, the estimates generally are aggregated for building groups and are reported at zip code levels (Lu et al. 2020). In addition, these methods require expert judgment and recalibration using historical data in order to be adapted to a specific event (Erdik et al. 2011). With the recent improvements in vision-based machine learning methods, remote sensing also has been studied extensively for postearthquake building damage detection (Kerle et al. 2020; Ji et al. 2020; Naito et al. 2020; Cooner et al. 2016). Although aerial images of the impacted region can be obtained shortly after the earthquake, environmental conditions such as cloud cover or haze can limit the applicability of these methods. Even with clear skies, these images can feature only the roof condition of a building, and may not identify internal cracks or non-structural damage. Therefore, the majority of the remote sensing studies consider a binary damage–no damage classification (Kerle et al. 2020).

To obtain a reliable damage assessment, on-ground inspection teams currently are deployed to the affected region for an accurate assessment of damage. These teams inspect individual buildings for damage and failure mechanisms, and report the results in a reconnaissance survey (Brando et al. 2017; Chiaro et al. 2015). In addition to damage estimation, valuable lessons can be learned from investigating the effects of an earthquake on different types of buildings, because this can help improve the seismic design codes. Currently, the labor-intensive reconnaissance surveys can cover only a limited number of buildings during the first few weeks after the earthquake. Moreover, the selection of buildings for damage inspection is based on the observable exterior damage. Not only is this process time-consuming and costly, but locating the damaged buildings can be a challenge itself, especially in the aftermath of an earthquake. Consequently, completing all damage inspections in the affected region can take several months or even years following the seismic event (Loos et al. 2020). Lenjani et al. (2020a) used fragility functions to estimate the damage state of buildings in order to prioritize the buildings for damage inspections.

¹Research Assistant, Dept. of Civil and Environmental Engineering, Univ. of Utah, Salt Lake City, UT 84112. ORCID: <https://orcid.org/0000-0003-2825-4937>. Email: m.sheibani@utah.edu

²Research Assistant, Dept. of Civil and Environmental Engineering, Univ. of Utah, Salt Lake City, UT 84112. Email: yinhu.wang@utah.edu

³Assistant Professor, Dept. of Civil and Environmental Engineering, Univ. of Utah, Salt Lake City, UT 84112 (corresponding author). ORCID: <https://orcid.org/0000-0001-7301-6765>. Email: ge.ou@utah.edu

⁴Assistant Professor, Dept. of Civil and Environmental Engineering, Univ. of Utah, Salt Lake City, UT 84112. Email: nikola.markovic@utah.edu

Note. This manuscript was submitted on September 30, 2020; approved on October 15, 2021; published online on December 10, 2021. Discussion period open until May 10, 2022; separate discussions must be submitted for individual papers. This paper is part of the *Journal of Engineering Mechanics*, © ASCE, ISSN 0733-9399.

Recently, data-driven methods have been suggested in different areas of earthquake engineering to enhance flexibility and reduce the computational burden in physics-based simulations (Xie et al. 2020; Hormozabad and Soto 2021). By treating various characteristics of the building structure and ground motion as inputs (or features), surrogate models are able to emulate the nonlinear response (output or labels) of the model without performing the expensive time-history analysis (de Lautour and Omenzetter 2009). For a regional postearthquake damage assessment, every inspected building provides an example of this input–output relationship and can be used to fit the surrogate model. Importantly, a relatively small number of collected damage observations can be used to train a model to predict accurately the damage to the uninspected buildings in the affected region (Sheibani and Qu 2020b). Such a data-driven method is adaptable to the impact of a specific event on the buildings constructed in the considered region. Furthermore, this approach will yield a high-resolution damage prediction map for individual buildings in the region, without relying on historical data or expert judgment, as do the conventional approximation methods.

As a powerful nonparametric method, Gaussian process regression (GPR) can be used to model the nonlinear relationship between the building features and damage. A properly trained GPR model is capable of predicting probabilistic building damage levels with promising accuracy (Sheibani and Ou 2020a). Results of the damage predictions for the entire building inventory in the affected region can be utilized further in the total loss estimation applications. With the benefit of the developed rapid and inclusive damage assessment method, an efficient method to implement the practical damage inspection scenarios was studied in this paper. In most supervised learning applications, training data are selected randomly, and the cost of data acquisition is either negligible or treated uniformly. However, for the regional damage assessment problem, a careless selection of buildings for inspection and training will lead to tremendous expense. The assessment of damage intensity for a building after the earthquake demands a considerable amount of time and resources. Therefore, an effective strategy is required to implement the postseismic building damage inspection in practice. Such an approach should maximize information gain given the constrained inspection resources (e.g., X crew members in Y teams working Z h/day). Thus, only the most informative or most representative buildings should be inspected. At the same time, the overlap in the information obtained from the survey should be minimized.

The main contribution of this paper is the development of a novel inspection scheduling framework for the regional postseismic

building damage inference. The framework integrates, in an innovative way, machine learning for damage inference and route optimization to achieve efficient data acquisition in a postearthquake reconnaissance survey. This is achieved as follows:

- Machine learning. The pool of available buildings in the seismic-impacted region is partitioned using the k -means clustering algorithm in the space of the building and earthquake features. The buildings assigned as the cluster centroids should be representative of other buildings in their clusters. Therefore, buildings that represent centroids of larger clusters are assigned higher inspection priority (i.e., higher expected reward from the inspection) because previous studies demonstrated that querying the labels for larger clusters increases the accuracy of regression algorithms when applied to collected data (Wu 2018).
- Route planning. To reach the desired damage inference accuracy as soon as possible, the order in which the candidate buildings are inspected should be determined while accounting for the available inspection resources (e.g., one team working 10 h/day). The problem of inspection scheduling, given the priority (i.e., reward) of each building and the available inspection resources, is formulated as an orienteering problem (OP). To our knowledge, this is the first application of the OP concerned with efficient information collection.
- Numerical evaluation. The effectiveness of the proposed framework (Fig. 1) is demonstrated on a simulated earthquake scenario in the San Francisco city area. The experimental results indicate that, within the same period, the proposed optimization method improves the accuracy of damage inference by 6% while reducing the total travel distance of the surveying team by 38%, compared with the situation in which inspection is conducted by simply prioritizing building with the larger weights (and without applying route optimization algorithms).

The remainder of this paper is organized as follows. Section “Regional Building Damage Prediction Model” briefly describes the theoretical foundations of the GPR, which is used as the surrogate model to infer buildings damage levels. The proposed inspection scheduling framework and the theory of the OP are presented in the section “Development of Strategies for Optimized Damage Inspection.” The advantages of using centroids of k -means clusters in the regional building damage inference problem is demonstrated on an earthquake testbed in the section “San Francisco Bay Area Regional Earthquake Simulation Testbed.” Section “Effect of Representative Sampling on Predictive Performance” investigates and compares the inference performance under different inspection and scheduling approaches, and the section “Conclusion” draws conclusions.




Fig. 1. Proposed framework for effective postearthquake damage assessment.

Regional Building Damage Prediction Model

In the proposed framework, the surrogate model is trained with a number of input–output examples. The inputs to the model are structural variables and ground motion intensity measures (IMs). Because explicit structural models, in particular mass, stiffness, and damping, are not accessible for real-world buildings without advanced testing, a set of implicit structural variables, including the number of stories, structural type, age, and so forth, is adopted. In addition, the IMs that have been shown to be correlated with the structural damage, such as spectral intensity, spectral acceleration, and so forth, are utilized to represent the input loading to the building. Sheibani and Ou (2020a) showed that an optimal learning can be performed using a set of six readily obtainable building variables and six earthquake characteristic features. On the other hand, various damage indexes such as the residual displacement, observable damage, economic loss, and so forth can be considered as the output to the model. A GPR model was trained on the sparse input–output examples to perform damage inferences for the unobserved buildings in the region.

Gaussian process (GP) models are nonparametric generalizations of the linear regression that make probabilistic inferences using Bayesian inference. A GP is a distribution over functions such that any finite number of points sampled at particular inputs from each function have a joint Gaussian distribution. This distribution is characterized by a mean and covariance function. Here, we point out only the final prediction equations; Williams and Rasmussen (2006) and Sheibani and Ou (2020b) provided a detailed derivation of the formulations and the model selection tasks for the regional damage assessment problem, respectively.

Assuming an m -dimensional training set with n datapoints $\{(\mathbf{x}_i, y_i) | i = 1, \dots, n\}$, and the set of all inputs $\mathcal{X}$, for any pair of input vectors $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}$, we have the output means $\mu(\mathbf{x}_1)$ and $\mu(\mathbf{x}_2)$, and the covariance $K(\mathbf{x}_1, \mathbf{x}_2)$. Given the observations $\mathbf{y}_{\mathcal{R}}$ for a subset $\mathcal{R} \subset \mathcal{X}$, the conditional probability distribution of label $y_{\mathbf{x}_*}$ for any new input $\mathbf{x}_* \in \mathcal{X} \setminus \mathcal{R}$ can be computed as $P(y_{\mathbf{x}_*} | \mathbf{y}_{\mathcal{R}})$ such that

$$\mu(\mathbf{x}_* | \mathcal{R}) = \Sigma_{\mathbf{x}_* \mathcal{R}} \Sigma_{\mathcal{R} \mathcal{R}}^{-1} \mathbf{y}_{\mathcal{R}} \quad (1)$$

$$\sigma^2(\mathbf{x}_* | \mathcal{R}) = K(\mathbf{x}_*, \mathbf{x}_*)^T - \Sigma_{\mathbf{x}_* \mathcal{R}} \Sigma_{\mathcal{R} \mathcal{R}}^{-1} \Sigma_{\mathcal{R} \mathbf{x}_*} \quad (2)$$

where $\Sigma_{\mathcal{R} \mathcal{R}}$ = covariance matrix for members of $\mathcal{R}$, every entry of which is calculated as $K(\dots)$. Although the mean function commonly is considered to be zero, the covariance function plays an intrinsic role in modeling the relations in the obtained experimental data. In accordance with the results of the model selection (Sheibani and Ou 2020a), the automatic relevance determination type of the rational quadratic (RQ) function was chosen for modeling the nonlinear behavior of data. For any two input vectors $\mathbf{x}_1$ and $\mathbf{x}_2$, the covariance can be computed as

$$K_{\text{RQ}}(\mathbf{x}_1, \mathbf{x}_2) = \sigma_f^2 (1 + (\mathbf{x}_1 - \mathbf{x}_2)^T \frac{\mathbf{M}}{2\alpha} (\mathbf{x}_1 - \mathbf{x}_2))^{-\alpha} \quad (3)$$

where σ_f^2 = signal variance; and $\mathbf{M} = \text{diag}(\mathbf{1})^{-2}$, where $\mathbf{1}$ = vector containing characteristic length scales, and $\alpha > 0$ determines the shape of the function. The mentioned parameters are tuned during the training procedure with maximum-likelihood estimation (MLE). The libraries of the GPML toolbox (Rasmussen and Nickisch 2010) were used for inference purposes in this paper.

Development of Strategies for Optimized Damage Inspection

As discussed previously, the optimization of inspection surveys has two objectives: (1) identifying the representative building candidates for inspection, and (2) optimizing the inspection scheduling and path given the weights of each candidate building. This section proposes using the k -means clustering method to identify the inspection candidates. For the inspection prioritization and determination of the most efficient route(s), several different scenarios are discussed, and the OP is formulated.

K-Means Clustering to Determine Weights

The predictive performance of the surrogate model improves with new datapoints (i.e., input–output examples). In our problem, these examples are in the form of building damage observations, which are expensive to make. Therefore, to reduce the data collection costs, it is beneficial to inspect a wide spectrum of buildings and reduce the amount of information overlap.

It has been shown that by clustering the input space, a subset of points can be obtained that is representatives of the data set (Wu 2018; Liu et al. 2020; Dehghani et al. 2020). Because the samples within a cluster are well correlated and expected to have similar outputs, the nearest datapoint to the center of the cluster (centroid) can be considered as the representative of that partition. These centroids are located far from each other in the feature space, and therefore we are meeting the diversity criterion in our data collection as well (Demir et al. 2010). The k -means clustering strategy can be used to project the input space into a subspace created by k representative points.

The k -means clustering algorithm is an iterative procedure that classifies a dataset into k disjoint clusters in which datapoints in each cluster are more similar to each other than to datapoints in other clusters. The Euclidean distance generally is used to determine the similarities between datapoints. This algorithm consists of two phases in each iteration. In Phase 1, k samples are chosen as centroids, and in Phase 2, the remaining samples are associated with the nearest centroid to shape k clusters. The next iteration starts with choosing the average value of all the points in each cluster as the new center, and Phase 2 is repeated similarly to the first iteration. This process terminates upon either convergence or reaching the maximum number of iterations. Further details of this algorithm were given by Nazeer and Sebastian (2009).

In this paper, the k -means clustering algorithm partitions the input feature space, and the centroids are considered as the most diverse and representative cases from the pool of available buildings in the region. Although k -means is less prone to outliers, we made sure to disregard clusters with only one member element.

Different Scenarios to Prioritize Inspections

In addition to choosing the most informative buildings for damage inspection, the order in which the inspections are made has a crucial impact on time efficiency. The centroids of the clusters that have larger sizes are representatives of more datapoints, and hence picking them earlier expedites the rate of reduction of the generalization error of GPR. The inspection routes should be efficient to enable more inspections in a limited amount of time (e.g., 10 work h/day). Therefore, in this paper, the objective is to both select most informative buildings and design efficient inspection routes for the surveying team.

Based on these considerations, we propose an efficient heuristic that meets both objectives. In this method, the informativeness of the inspected buildings is ensured by selecting candidates from the

Table 1. Different damage inspection scenarios considered

Scenario	Objective	Description
S1	Shortest time	The nearest buildings to the depot from the pool of all buildings are chosen for inspection
S2	Fastest convergence	Building candidates are chosen as cluster centroids with the priority given to larger clusters in the inspection queue
S3	Shortest route	Candidate buildings are chosen as cluster centroids with a queue set for the maximum number of inspections in each day
S4	Fastest convergence with route optimization	Candidate buildings are chosen as cluster centroids. Maximum number of inspections in each day is sought with priority given to the centroids of the larger clusters

centroids of clusters. We also prioritize the centroids of larger clusters in order to increase the rate of convergence. Furthermore, to optimize the surveying costs and time, we use OP to assign an efficient route for daily tasks of the inspection team. The proposed inspection method (S4) and three other scenarios (S1, S2, and S3) were studied (Table 1). S1–S3 are alternative approaches that can be pursued to fulfill the datapoint selection parts of the framework (Fig. 1). For each scenario, a fixed depot location is considered at which the surveying team starts and ends the daily inspection routine.

The problem of prioritizing buildings for inspection and determining the most efficient route(s) in Scenarios S3 and S4 are formulated as the OP, which is concerned with maximizing the reward collected from visiting nodes given the allowed time constraint. In the approach taken for S4, the cluster sizes are considered as reward amounts. These rewards are identical in S3 to allow for the best route, regardless of the reward amount. The subsequent sections explain details of the OP.

Orienteering Problem

The OP is a well-studied combinatorial optimization problem. In this problem, a set of nodes is given, and a score is associated with each node. The goal of the OP is to determine a time-constrained path to visit a subset of these nodes such that the total reward collected is maximized. The name of the OP is derived from a sports game of orienteering (Chao et al. 1996); the problem is also called the selective traveling salesman problem, the maximum collection problem, and the bank robber problem (Vansteenwegen et al. 2011). A large variety of practical problems, such as home fuel delivery (Golden et al. 1987), tourist trips (Yu et al. 2015), and route design for special events such as wildfires (van der Merwe et al. 2014) can be modeled as an OP. In this paper, we propose the first application of the OP in the context of efficient postseismic reconnaissance surveys, in which the team is supposed to collect as much information via building inspection in a limited amount of time.

Problem Formulation

For completeness, we present the mathematical programming formulation of the OP used to design the most efficient routes for the survey team. We proceed with the so-called two-index formulation, which is defined on a graph $G(V, A)$, where V denotes a set of k vertices (i.e., candidate buildings for inspection determined as centroids as part of k -means clustering) and A is a set of arcs; S_i is defined as the weight (or reward) for visiting vertex $i \in V$. The survey team starts its trip from a depot denoted 0 and returns to the depot after finishing its daily inspection task. Let r_{ij} be a binary variable, which takes the value of 1 if the arc (i, j) is included in the team's route, and 0 otherwise. The travel time from node i to node j is denoted t_{ij} , and the total travel time cannot exceed a given time restriction $T_{\max}$. Lastly, u_i is an integer variable representing

the position of node i in the path. The math programming formulation (Vansteenwegen et al. 2011)

$$\max_{r_{ij} \in \{0,1\}, u_i \in \mathbb{Z}^+} \sum_{i \in V \setminus \{0\}} \sum_{j \in V} S_i r_{ij} \quad (4)$$

$$\sum_{j \in V \setminus \{0\}} r_{0j} = \sum_{i \in V \setminus \{0\}} r_{i0} = 1 \quad (5)$$

$$\sum_{i \in V} r_{ih} = \sum_{j \in V} r_{hj} \leq 1 \quad \forall h \in V \setminus \{0\} \quad (6)$$

$$\sum_{i \in V} \sum_{j \in V} t_{ij} r_{ij} \leq T_{\max} \quad (7)$$

$$1 \leq u_i \leq k \quad \forall i \in V \quad (8)$$

$$u_i - u_j + 1 \leq (k-1)(1 - r_{ij}) \quad \forall i, j \in V \quad (9)$$

seeks to maximize the total collected weights. Eq. (5) ensures that the route of survey team starts and ends at the depot. Eq. (6) guarantees that the route is complete and that each intermediate node in the route is visited at most once. Eq. (7) ensures that the route must satisfy a maximum travel duration $T_{\max}$. Eqs. (8) and (9) prevent subtours.

OP Heuristics

The OP is a notoriously difficult combinatorial optimization problem, and it cannot be solved optimally for instances involving several hundred nodes. As a result, many heuristics have been developed to tackle large instances of the OP, such as evolutionary algorithms and variable neighborhood search methods (Verbeeck et al. 2014; Kobeaga et al. 2018; Palomo-Martínez et al. 2017), which typically provide near-optimal solutions. In our framework, we apply a combination of an insertion heuristic and a 2-opt algorithm for route construction and improvement. The proposed approach sorts nodes based on their weight (or reward) and myopically inserts unvisited nodes with the highest reward until the constraint on the total travel time $T_{\max}$ is violated. This initial route then is made more efficient by applying the 2-opt algorithm, which changes the order of nodes to avoid having the route cross itself (Engels and Manthey 2009). Specifically, it makes improvements by exchanging two edges of the previously constructed tour with two other edges until a local optimum is reached. The OP heuristics are applied daily until enough buildings have been inspected (i.e., until the GPR model reaches the desired accuracy). The details of the adopted approach are presented in Algorithm 1. Insertion algorithms are applied widely due to their effectiveness on large-scale routing problems (Bertsimas et al. 2019), and the 2-opt local search heuristic commonly is applied to improve routes further. Fig. 2 presents an illustrative example of a route obtained using the insertion algorithm and further improved using the 2-opt heuristic. After implementing the improvement

procedure (Kay 2013), the total travel cost decreased from 108 to 102 km.

Algorithm 1. Insertion and 2-opt heuristic

Data: initial tour $P = \emptyset$, travel time t_{ij} from node i to j , cumulative travel time of tour T_P , travel time constraint $T_{\max}$, k nodes with descending weights, depot v and the last node ν in tour P , and working time ξ

Result: a tour P maximizing collected weights given the maximum allowed travel time $T_{\max}$

$P \leftarrow P \cup \{v\};$

$T_P \leftarrow 0;$

for $i = 1:N$ **do**

if node i is unvisited **then**

 travel time from ν to i : $t_{\nu,i};$

 travel time from ν to depot: $t_{i,v};$

If $T_P + t_{\nu,i} + \xi + t_{i,v} \leq T_{\max}$ **then**

$P \leftarrow P \cup \{i\};$

 apply 2-opt algorithm and return an improved tour P^* ;

$P \leftarrow P^*;$

$T_P \leftarrow T_{P^*};$

end if

end if

end for

San Francisco Bay Area Regional Earthquake Simulation Testbed

The methodologies proposed in this paper were evaluated in a standard simulation testbed which emulated the regional seismic impact and structural performance under physical-based modeling. The detailed description of the earthquake simulation methods and the data obtained from the simulation are discussed in this section.

A hypothetical earthquake with a rupture along the Hayward Fault was considered, which produced ground motions with PGAs as high as $0.64g$ in the San Francisco city area. This scenario was simulated using the open-source modular Regional Workflow for Hazard And Loss Estimation (rWHALE version 1.1.0) provided

by the Natural Hazards Engineering Research Infrastructure (NHERI) SimCenter (Lu et al. 2020). Basic information of 20,000 buildings distributed in the San Francisco city area was used to create multiple degree of freedom (MDOF) finite-element models for this study. A grid of ground motions was determined in the region using the SW4 wave propagation method (Rodgers et al. 2019), and the nearest ground motion was assigned to each building [Fig. 3(a)]. The simulation was developed based on a nonlinear MDOF model for every building in the region, and different damage indexes were calculated using time-history analysis. Due to the use of actual building information and the immense consideration of uncertainty in the rWHALE program, the results of the simulations were considered to be close representations of damage labels for a real-world earthquake scenario. Details of the data generation process were given by Lu et al. (2020) and Sheibani and Ou (2020a). In this paper, a building's economic loss ratio is considered as the damage label for inference. The geographical distribution of the true labels is shown in Fig. 3(b). Because the quantities of this label are bounded between 0 and 1, labels were transformed with the inverse cumulative Gaussian function (probit) to satisfy the Gaussian noise assumption in the GPR formulations, and then transformed back to the label space for performance assessment.

Results

The results demonstrate the efficient building damage inspection solution for postseismic building damage inference according to the two objectives: (1) the representativeness of the inspection candidates identified using k -means clustering, and (2) the impact of inspection scheduling optimization to the inference efficiency.

Effect of Representative Sampling on Predictive Performance

As mentioned in the section "Development of Strategies for Optimized Damage Inspection," the hypothesis is that the improvement in predictive performance should increase if informative and diverse datapoints are selected for training. To validate the contribution of representative sampling on damage inference accuracy, the predictive performance with k -clustered sampling was compared




Fig. 2. Illustrative example of vehicle routing derived with insertion algorithm and route improvement using 2-opt algorithm: (a) route based on the insertion algorithm; and (b) route improvement using 2-opt heuristic. (Map data © 2020 Google.)





Fig. 3. Distribution of data on San Francisco map: (a) ground motion locations and their corresponding PGA; and (b) calculated economic loss ratio for each building.

with a random sampling approach. It was assumed that the budget was available for a maximum of 300 building damage inspections, and the cost of inspecting each building was treated as uniform at this point. Therefore, the pool of all buildings was partitioned into 300 clusters, and the centroid of each cluster was used as the representative for that cluster. The predictive performance was evaluated using the relative difference (RD) error as

$$\text{RD}(\mu_{\mathbf{x}_*}, y_{\mathbf{x}_*}) = \frac{2(\mu_{\mathbf{x}_*} - y_{\mathbf{x}_*})}{|\mu_{\mathbf{x}_*}| + |y_{\mathbf{x}_*}|} \quad (10)$$

Using this formula, RD is always between 0 and 2.

Fig. 4 shows the result of the comparison averaged over 100 realizations. A higher percentage of the testing points was predicted within the error margins specified when sampling was based on clustering. The difference was more pronounced when the error

margin was at a lower level. For example, Fig. 4(a) indicates that to achieve predictions with $RD < 0.1$ for about 56% of datapoints, 150 training points are required when sampling is based on clustering, whereas this number should be increased to 300 points when sampling is performed randomly. Therefore, using clustering to choose candidate buildings for damage inspection reduces the total number of inspections required.

Moreover, it is critical to observe the effectiveness of adding the representative of a cluster to the training set on the prediction accuracy of the points within the same cluster. Therefore, we ranked clusters from large to small and grouped them into 10 groups. To refer to these cluster groups more easily, each group was named using the convention CX, where X is an integer from 1 to 10. Therefore, Group C1 contained the 30 largest clusters, whereas C10 included the smallest clusters in the data. To display the prediction accuracy of the group members before and after their





Fig. 4. Comparison of the prediction accuracy for random selection of training data and clustering selection, considering different training set sizes: (a) $RD < 0.1$; and (b) $RD < 0.2$.




Fig. 5. Improvement in prediction accuracy of different cluster groups after the representatives of each group are inspected.




Fig. 6. Influence of the inspection of the representatives of each cluster group on the accuracy of all datapoints along with the number of members in each group. Error bars show the first and third quartiles.

representatives were inspected, the set of 30 centroids in each group was inspected and the prediction accuracy of the members of the group was evaluated. It also is advantageous to see the improvements in prediction accuracy of the members in other groups. Therefore, all combinations were evaluated (Fig. 5). The maximum improvement in accuracy for the members of a group is obtained when their representatives were inspected exclusively. The improvements are less significant for the members of other groups. However, the members of other groups benefit significantly more when representatives of larger clusters are inspected. To evaluate the results more quantitatively, the percentage of all datapoints for which the accuracy moved into the error margin when the representatives of each group were inspected is shown in Fig. 6. The total number of members within each group relative to the total number of datapoints also displayed in Fig. 6. Inspecting the representatives of larger clusters influence the prediction accuracy of more datapoints. For example, inspecting the 30 representatives of C1 moved the prediction accuracy of about 60% of all datapoints to the error margin, whereas this percentage was below 5% for C10.

Therefore, it can be stated that starting with larger clusters in the damage inspection scenarios significantly can increase the rate of

improvement in the generalization error and reduce the time required to reach higher prediction accuracy. These observations demonstrate that (1) using the k -means clustering sampling, one can identify a representative datapoint for each cluster by selecting the cluster centroid; and (2) it is advantageous to start the inspection procedure from the representatives (training points) of larger clusters to cover a larger number of electors (testing points) during the first days of inspection.

Comparison of Inference Efficiency under Different Inspection Scheduling Scenarios

To fulfill the second objective in developing a practical postseismic inspection practice, the impact of the inspections scheduling and route optimization on the reconnaissance efficiency and costs was evaluated. According to ATC-38, a building inspection takes about 2 person-h (ATC 2000). Thus, to compare the required time for each approach, the effective work hours spent by a team of four inspectors was calculated at 0.5 h/building. A maximum of 10 h/day of work budget for the inspection team was considered, which included the travel times to and from the depot. As an approximation, the travel path between two nodes was considered to be equal to the great circle distance between the geographical locations of the buildings, which was obtained using Haversine formula. Three different average travel speeds (10, 25, and 40 km/h) were considered to account for different difficulties in the aftermath of the earthquake, such as probable road and bridge closures. Before reaching any conclusion from results, it should be noted that there are two independent variables that can vary the outcome of the inference: the random initial seed selected for k -means clustering, and the depot position assumed for route planning. To illustrate the detailed dynamics due to each inspection scheduling method, as well as the deterministic patterns emerged from the two random factors, the inspection scheduling performance was investigated using a single realization first, and then a statistical summary of a batch Monte Carlo realization.

Performance Evaluation of Single Inspection Scheduling Realizations

To illustrate the improvements of the inference performance to the scheduling dynamics, a single realization was considered. The depot location was assumed to be the same across four scenarios, and an average travel speed of 10 km/h was considered.

The sample paths from the proposed four scheduling scenarios are depicted in Fig. 7. To reiterate, the objectives were the overall shortest time for S1, prioritizing larger clusters for S2, inspecting as many representatives as possible in each day for S3, and gathering the most reward in each day by solving the OP for S4. Although several days were needed to complete the damage inspection for all candidate buildings, Fig. 7 demonstrates only the paths of the first 4 days for better clarity. The objectives heavily impacted the inspection paths recommended for the surveying team. The paths recommended for S1 were based solely on the proximity of buildings to the depot location, and hence, the travel distances were insignificant. The routes suggested for S3 carefully implemented the traveling salesman heuristics and provided optimum paths that visited as many candidates as possible in each day, regardless of the node weights. Paths recommended for S2 and S4 prioritized the nodes based on their weights. However, because no path optimization method was performed in S2, routes crossed themselves and sub-optimal traveling behavior was proposed for the team. With the implementation of the OP in S4, routes were planned carefully to prioritize nodes with higher weights while reducing traveling times in order to allow the team to visit as many nodes as possible.




Fig. 7. Sample inspection paths taken by the team during the first 4 days for different inspection scenarios. In figures illustrating Scenarios S2, S3, and S4, the marker sizes illustrate the relative cluster sizes.

Table 2. Details of inspection surveys after first 2 weeks following earthquake occurrence

Scenario	Number of inspections performed	Percentage of predictions with $RD < 0.20$	Total distance traveled (km)	Total travel time (min)
S1	190	29.7	17.3	104
S2	93	73.3	483.0	2,897
S3	162	67.7	170.4	1,022
S4	135	77.4	297.4	1,784

The two important factors in the damage estimation process are prediction accuracy and the required time to achieve the desired level of accuracy. Therefore, Table 2 compares the sample inspection scenarios based on these factors. Considering 5 workdays/week, after the first two crucial weeks following the occurrence of the earthquake, by pursuing the S4 scenario, the GPR algorithm predicted the damage labels of 77.4% of the buildings with $RD < 0.2$.

Compared with the S2 scenario, the surveying in S4 required significantly less travel in the region yet was able to inspect more buildings.

Statistical Inference Summary to Evaluate Different Inspection Scenarios

To investigate the deterministic pattern of the four inspection scenarios which are independent of the depot location or clustering seeds, 100 Monte Carlo realizations were performed for each scenario with different initial seeds for the k -means clustering (except for S1) and also with a random location for the depot. However, in the four scenarios, the initial seed and depot location for each realization are kept the same. The first three quartiles of daily improvements in the performance of GPR are shown in Fig. 8 for the first 3 weeks of work. The highest rate of daily improvement in predictive performance was for the S4 inspection scenario. For example, considering the case of 10 km/h average travel speed [Fig. 8(a)], after only 5 days of inspection the labels were predicted with $RD < 0.2$ for 71% of the datapoints with Scenario S4, whereas this




Fig. 8. Daily improvement in predictive performance for different inspection scenarios. Error bars show the first and third quartiles. Inspected time (day): (a) travel speed = 10 km/h; (b) travel speed = 25 km/h; and (c) travel speed = 40 km/h.




Fig. 9. Cumulative distance travelled for different inspection scenarios versus improvement of predictive performance. Error bars show the first and third quartiles. Cumulative distance traveled: (a) travel speed = 10 km/h; (b) travel speed = 25 km/h; and (c) travel speed = 40 km/h.

percentage was 65% and 55% with Scenarios S2 and S3, respectively. At 25 km/h average speed, about 9 days was required to bring the accuracy to 80% with S4, whereas 11 and 15 days were required for S2 and S3, respectively. Although the overall performance increased for S2, S3, and S4 at 40 km/h, the differences between the scenarios remained almost the same. Furthermore, the predictive performance never reached acceptable levels when representative points were not chosen for inspection in the S1 inspection scenario.

To compare the overall distance traveled by the inspection team, the cumulative distances traveled for every scenario is shown in Fig. 9. As in Fig. 8, the results are shown for the first 3 weeks of inspection. Scenarios S2 and S4 in this figure demonstrate the considerable reduction in travel distance by solving the OP for the inspection route optimization. The surveying team in S4 was able to inspect more buildings in a limited time, yet traveled almost 35% less than that in S2 at 10 km/h. The savings in total travel distance was about 51% and 56% at 25 and 40 km/h, respectively, in S4 compared with S2.

The efficiency of the inspection scenarios also can be compared by computing the ratio of the inspection hours and travel hours. Fig. 10 shows the individual amounts of time that the team spent traveling versus the effective inspection hours. As expected, in the

S1 scenario, the team spent the least amount of time on travel, which makes this scenario efficient from this point of view. However, inspections made in this scenario did not improve the predictive performance of the GPR algorithm (Figs. 8 and 9). The team also spent relatively short amounts of time traveling in the S3 scenario, in which the objective was the minimum travel distance regardless of the cluster sizes. Although the ratio of traveling hours was reduced, because the rate of improvement in the predictive performance was not optimal, inspections made with the S3 scenario were not favorable for rapid damage assessment purposes. Comparing S2 and S4 demonstrates the improvement in the ratio of inspection hours to travel hours. At the 10 km/h speed, at the end of the third week, the team had spent almost 51% of their time traveling between buildings with S2, whereas this percentage decreased to 32% with S4.

Finally, it is beneficial to see the improvement in the predictive performance of the GPR for the members of each cluster after their representative is inspected. For better visualization, all 300 clusters were divided into 10 groups based on their sizes, and the daily accuracy of prediction for their members was evaluated using the RD criterion. Fig. 11 shows the percentage of members of each cluster group that were predicted with $RD < 0.2$ before and after the representatives of the clusters were inspected. Fig. 11 considered only




Fig. 10. Cumulative weekly work hours for different inspection scenarios: (a) travel speed = 10 km/h; (b) travel speed = 25 km/h; and (c) travel speed = 40 km/h.




Fig. 11. The progress of the prediction accuracy for the members of different clusters: (a) S2; (b) S3; and (c) S4.

the 10 km/h average speed, and the results were averaged over 100 realizations. The prediction accuracy lines have sharp transitions at the time when the representatives are inspected. It can be inferred that using the S4 method, except for the two largest cluster groups for which the performance of S2 was comparable, all other cluster groups reached high predictive performance in a shorter period [Figs. 11(a and c)]. Because all cluster groups had the same priority in the inspection queue, Fig. 11(b) shows a rather uniform distribution for the average day in which the representative of each group was inspected. Due to this fact, the predictive performance stays at lower levels during the first days of the inspection scenario.

Discussion

Considering all aspects of comparison between different scenarios, it is clear that the proposed method significantly improves the damage assessment efficiency by leveraging the OP to plan the team's daily tasks. This scenario can be used to guide the reconnaissance surveying teams systematically and rapidly provide a region-wide damage estimation after seismic events. Not only does it reduce the inspection costs and time, but the availability of the PDNA within a shorter time expedites the recovery efforts. Finally, although the method was validated on a simulated earthquake testbed, it is applicable to a real postearthquake damage inspection scenario. However, other

constraints also should be studied for a real scenario, such as road closures, multiple inspection teams working concurrently, unavailability of building data in some regions, and so forth.

Conclusion

An efficient data-based framework for inferring structural damage after earthquakes is proposed. In the proposed method, the inventory of basic building information, available ground motion characteristics at buildings' sites, and sparse field observations after the event are used to train a regional building damage inference model. Gaussian process regression was selected to provide a high-performance emulation with small training sets. It was shown that the proposed inspection scheduling approach efficiently guides the surveying team in their daily tasks by recommending the most informative buildings for inspection and designing efficient inspection routes. In a comparison with other approaches, it was concluded that, without careful planning, it takes significantly longer for the prediction accuracy of the surrogate model to reach acceptable levels. In fact, if a surveying team inspects buildings at random, which usually is the case in practice, the model can predict accurately the damage for about 30% of buildings based on 2 weeks of inspections, whereas this percentage can reach about 77% if the proposed framework is followed.

Data Availability Statement

All data, models, or code that support the findings of this study are available from the corresponding author upon reasonable request.

Acknowledgments

This material is based upon work supported by the University of Utah and the National Science Foundation under Award Nos. 2004658 and 2112758. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

- ATC (Applied Technology Council). 2000. *Database on the performance of structures near strong-motion recordings: 1994 Northridge, California, earthquake*. Redwood City, CA: ATC.
- Bertsimas, D., P. Jaillet, and S. Martin. 2019. "Online vehicle routing: The edge of optimization in large-scale applications." *Oper. Res.* 67 (1): 143–162. <https://doi.org/10.1287/opre.2018.1763>.
- Bland, H. M., and J. D. Frost. 2013. "Opportunities and considerations for smartphone applications and mobile social media in post extreme event reconnaissance data collection." In *Forensic engineering 2012: Gateway to a safer tomorrow*, 505–514. Reston, VA: ASCE.
- Brando, G., et al. 2017. "Damage reconnaissance of unreinforced masonry bearing wall buildings after the 2015 Gorkha, Nepal, earthquake." *Earthquake Spectra* 33 (1): 243–273. <https://doi.org/10.1193/010817eqs009m>.
- Chao, I.-M., B. L. Golden, and E. A. Wasil. 1996. "A fast and effective heuristic for the orienteering problem." *Eur. J. Oper. Res.* 88 (3): 475–489. [https://doi.org/10.1016/0377-2217\(95\)00035-6](https://doi.org/10.1016/0377-2217(95)00035-6).
- Chiari, G., T. Kiyota, R. M. Pokhrel, K. Goda, T. Katagiri, and K. Sharma. 2015. "Reconnaissance report on geotechnical and structural damage caused by the 2015 Gorkha Earthquake, Nepal." *Soils Found.* 55 (5): 1030–1043. <https://doi.org/10.1016/j.sandf.2015.09.006>.
- Cooner, A. J., Y. Shao, and J. B. Campbell. 2016. "Detection of urban damage using remote sensing and machine learning algorithms: Revisiting the 2010 Haiti earthquake." *Remote Sens.* 8 (10): 868. <https://doi.org/10.3390/rs8100868>.
- Dehghani, N. L., Y. M. Darestani, and A. Shafieezadeh. 2020. "Optimal life-cycle resilience enhancement of aging power distribution systems: A MINLP-based preventive maintenance planning." *IEEE Access* 8 (Jan): 22324–22334. <https://doi.org/10.1109/ACCESS.2020.2969997>.
- de Lautour, O. R., and P. Omenzetter. 2009. "Prediction of seismic-induced structural damage using artificial neural networks." *Eng. Struct.* 31 (2): 600–606. <https://doi.org/10.1016/j.engstruct.2008.11.010>.
- Demir, B., C. Persello, and L. Bruzzone. 2010. "Batch-mode active-learning methods for the interactive classification of remote sensing images." *IEEE Trans. Geosci. Remote Sens.* 49 (3): 1014–1031. <https://doi.org/10.1109/TGRS.2010.2072929>.
- Engels, C., and B. Manthey. 2009. "Average-case approximation ratio of the 2-opt algorithm for the TSP." *Oper. Res. Lett.* 37 (2): 83–84. <https://doi.org/10.1016/j.orl.2008.12.002>.
- Erdik, M., K. Şeşetyan, M. Demircioğlu, U. Hancilar, and C. Zülfiyar. 2011. "Rapid earthquake loss assessment after damaging earthquakes." *Soil Dyn. Earthquake Eng.* 31 (2): 247–266. <https://doi.org/10.1016/j.soildyn.2010.03.009>.
- Golden, B. L., L. Levy, and R. Vohra. 1987. "The orienteering problem." *Naval Res. Logist.* 34 (3): 307–318. [https://doi.org/10.1002/1520-6750\(198706\)34:3<307::AID-NAV3220340302>3.0.CO;2-D](https://doi.org/10.1002/1520-6750(198706)34:3<307::AID-NAV3220340302>3.0.CO;2-D).
- Hormozabad, S. J., and M. G. Soto. 2021. "Load balancing and neural dynamic model to optimize replicator dynamics controllers for vibration reduction of highway bridge structures." *Eng. Appl. Artif. Intell.* 99 (Mar): 104138. <https://doi.org/10.1016/j.engappai.2020.104138>.
- Ji, M., L. Liu, R. Zhang, and F. Buchroithner. 2020. "Discrimination of earthquake-induced building destruction from space using a pretrained CNN model." *Appl. Sci.* 10 (2): 602. <https://doi.org/10.3390/app10020602>.
- Kay, M. G. 2013. "Matlog: Logistics engineering MATLAB toolbox." Accessed July 26, 2020. <https://people.engr.ncsu.edu/kay/matlog/MatlogRef.htm>.
- Kerle, N., F. Nex, M. Gerke, D. Duarte, and A. Vetrivel. 2020. "UAV-based structural damage mapping: A review." *ISPRS Int. J. Geo-Inf.* 9 (1): 14. <https://doi.org/10.3390/ijgi9010014>.
- Kobeaga, G., M. Merino, and J. A. Lozano. 2018. "An efficient evolutionary algorithm for the orienteering problem." *Comput. Oper. Res.* 90 (Feb): 42–59. <https://doi.org/10.1016/j.cor.2017.09.003>.
- Lenjani, A., I. Billionis, S. J. Dyke, and R. Monteiro. 2020a. "A resilience-based method for prioritizing post-event building inspections." *Nat. Hazard.* 100 (2): 877–896. <https://doi.org/10.1007/s11069-019-03849-0>.
- Lenjani, A., C. M. Yeum, S. Dyke, and I. Billionis. 2020b. "Automated building image extraction from 360° panoramas for postdisaster evaluation." *Comput.-Aided Civ. Infrastruct. Eng.* 35 (3): 241–257. <https://doi.org/10.1111/mice.12493>.
- Liu, H., Y.-S. Ong, X. Shen, and J. Cai. 2020. "When Gaussian process meets big data: A review of scalable GPs." *IEEE Trans. Neural Networks Learn. Syst.* 31 (11): 4405–4423. <https://doi.org/10.1109/TNNLS.2019.2957109>.
- Loos, S., D. Lallemand, J. Bakera, J. McCaughey, S.-H. Yunc, N. Budhathokid, F. Khanb, and R. Singhd. 2020. "G-DIF: A geospatial data integration framework to rapidly estimate post-earthquake damage." *Earthquake Spectra* 36 (4): 1695–1718. <https://doi.org/10.1177/8755293020926190>.
- Lu, X., F. McKenna, Q. Cheng, Z. Xu, X. Zeng, and S. A. Mahin. 2020. "An open-source framework for regional earthquake loss estimation using the city-scale nonlinear time history analysis." *Earthquake Spectra* 36 (2): 806–831. <https://doi.org/10.1177/8755293019891724>.
- Naito, S., H. Tomozawa, Y. Mori, T. Nagata, N. Monma, H. Nakamura, H. Fujiwara, and G. Shoji. 2020. "Building-damage detection method based on machine learning utilizing aerial photographs of the Kumamoto earthquake." *Earthquake Spectra* 36 (3): 1166–1187. <https://doi.org/10.1177/8755293019901309>.
- Nazeer, K. A., and M. Sebastian. 2009. "Improving the accuracy and efficiency of the k-means clustering algorithm." In *Proc., World Congress on Engineering*, 1–3. London: Association of Engineers.
- Palomo-Martinez, P. J., M. A. Salazar-Aguilar, G. Laporte, and A. Langevin. 2017. "A hybrid variable neighborhood search for the orienteering problem with mandatory visits and exclusionary constraints." *Comput. Oper. Res.* 78 (Feb): 408–419. <https://doi.org/10.1016/j.cor.2015.11.007>.
- Rasmussen, C. E., and H. Nickisch. 2010. "Gaussian processes for machine learning (GPML) toolbox." *J. Mach. Learn. Res.* 11 (Dec): 3011–3015.
- Rodgers, A. J., N. A. Petersson, A. Pitarka, D. B. McCallen, B. Sjogreen, and N. Abrahamson. 2019. "Broadband (0–5 Hz) fully deterministic 3D ground-motion simulations of a magnitude 7.0 Hayward Fault earthquake: Comparison with empirical ground-motion models and 3D path and site effects from source normalized intensities." *Seismol. Res. Lett.* 90 (3): 1268–1284. <https://doi.org/10.1785/0220180261>.
- Sheibani, M., and G. Ou. 2020a. "The development of Gaussian process regression for effective regional post-earthquake building damage inference." *Comput.-Aided Civ. Infrastruct. Eng.* 36 (3): 264–288. <https://doi.org/10.1111/mice.12630>.
- Sheibani, M., and G. Ou. 2020b. "Effective learning of post-seismic building damage with sparse observations." In *Vol. 3 of Model validation and uncertainty quantification*, 365–373. Cham, Switzerland: Springer.
- Steelman, J. S., and J. F. Hajjar. 2009. "Influence of inelastic seismic response modeling on regional loss estimation." *Eng. Struct.* 31 (12): 2976–2987. <https://doi.org/10.1016/j.engstruct.2009.07.026>.
- Ural, S., E. Hussain, K. Kim, C.-S. Fu, and J. Shan. 2011. "Building extraction and rubble mapping for city Port-au-Prince post-2010 earthquake with GeoEye-1 imagery and lidar data." *Photogramm. Eng. Remote Sens.* 77 (10): 1011–1023. <https://doi.org/10.14358/PERS.77.10.1011>.
- van der Merwe, M., J. Minas, M. Ozlen, and J. Hearne. 2014. "The cooperative orienteering problem with time windows." *Optim. Online* 2014 (7): 11.
- Vansteenkoven, P., W. Souffriau, and D. Van Oudheusden. 2011. "The orienteering problem: A survey." *Eur. J. Oper. Res.* 209 (1): 1–10. <https://doi.org/10.1016/j.ejor.2010.03.045>.

- Verbeeck, C., K. Sørensen, E.-H. Aghezzaf, and P. Vansteenwegen. 2014. "A fast solution method for the time-dependent orienteering problem." *Eur. J. Oper. Res.* 236 (2): 419–432. <https://doi.org/10.1016/j.ejor.2013.11.038>.
- Wilkinson, S., H. Stone, D. D'ayala, E. Verrucci, P. James, T. Rossetto, E. So, and C. Ellul. 2018. "How can new technologies help us with earthquake reconnaissance?" In *Proc., 11th US National Conf. on Earthquake Engineering*. Tyne, England: Newcastle Univ.
- Williams, C. K., and C. E. Rasmussen. 2006. *Gaussian processes for machine learning*. Cambridge, MA: MIT Press.
- Wu, D. 2018. "Pool-based sequential active learning for regression." *IEEE Trans. Neural Networks Learn. Syst.* 30 (5): 1348–1359. <https://doi.org/10.1109/TNNLS.2018.2868649>.
- Xie, Y., M. Ebad Sichani, J. E. Padgett, and R. DesRoches. 2020. "The promise of implementing machine learning in earthquake engineering: A state-of-the-art review." *Earthquake Spectra* 36 (4): 1769–1801. <https://doi.org/10.1177/8755293020919419>.
- Xin, D., J. E. Daniell, and F. Wenzel. 2018. "State of the art of fragility analysis for major building types in China with implications for intensity-PGA relationships." *Nat. Hazards Earth Syst. Sci.* 2018 (Sep): 1–34. <https://doi.org/10.5194/nhess-2018-254>.
- Yu, J., J. Aslam, S. Karaman, and D. Rus. 2015. "Anytime planning of optimal schedules for a mobile sensing robot." In *Proc., 2015 IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 5279–5286. New York: IEEE.