18TH WORLD CONFERENCE ON EARTHQUAKE ENGINEERING
WCEE2024
MILAN

MILAN, ITALY
30TH JUNE - 5TH JULY 2024

www.wcee2024.it

MACHINE LEARNING-BASED OPTIMIZATION FRAMEWORK TO SUPPORT RECOVERY-BASED DESIGN

O. Issa[1], R. Silva-Lopez[2], J.W. Baker[2], H.V. Burton[3]

[1] Department of Civil and Environmental Engineering, Stanford University, U.S.A., oissa@stanford.edu

[2] Department of Civil and Environmental Engineering, Stanford University, U.S.A.

[3] Department of Civil and Environmental Engineering, University of California, Los Angeles, U.S.A.

Abstract: *A study by FEMA suggests that 20-40% of modern code-conforming buildings in the United States would be unfit for re-occupancy following a major earthquake (taking months or years to repair) and 15-20% would be rendered irreparable. To address this gap, recovery-based design has been proposed to link building-level engineering with broader community resilience goals. However, the relationship between above-code design improvements and recovery performance is highly nonlinear and varies on a building-and site-specific basis, presenting a challenge to engineers, code developers, and policymakers. Furthermore, downtime simulations (e.g., FEMA P-58 + ATC-138) are computationally expensive, hindering exploration of the full design space. This paper summarizes our recent efforts to develop a framework to rapidly identify optimal above-code design improvements to achieve building-specific recovery objectives. The framework leverages surrogate models to allow for computationally efficient probing of recovery performance under a range of user-defined improvements, and optimization techniques that can be repeated for different stakeholder priorities. We demonstrate the framework through analysis of two modern steel building archetypes. For each building, optimal structural and nonstructural component improvements are identified for a range of recovery performance objectives. The study illustrates how the framework can be applied at scale to support the selection of building-specific recovery performance objectives.*

1   Introduction

The increasing human and economic exposure in seismically active regions underscores the need to bridge the gap between national seismic design provisions (which do not consider time to recovery) and community resilience goals. Using current design provisions, many at-risk communities will struggle to meet recovery time goals (e.g., SPUR, 2009), and control economic loss, with recent estimates placing the national expected annual loss from earthquakes at $14.7 billion (FEMA-USGS, 2023).

To address this issue, functional recovery has been proposed as a building performance objective that explicitly links design with organizational- or community-level resilience goals (EERI, 2019). Buildings designed for functional recovery are expected to recover their basic, tenant-specific functions within a target time, T_{target}, and also satisfy existing life safety objectives. Research in the area of performance-based earthquake engineering (PBEE), coupled with the emergence of enabling software (e.g., Zsarnóczay and Deierlein, 2020, HB-Risk, 2023) has led to early adoptions of recovery-based design by individual owners (Zimmerman and Herdrich, 2022; Mar and Aher, 2022; Forell/Elsesser Engineers, Inc., 2018). However, current applications of recovery-based design are typically developed through trial-and-error approaches. These approaches are tedious, since a variety of design strategies can be used to achieve a given post-earthquake recovery performance objective. Furthermore, the relationship between these improvements and recovery performance is highly nonlinear. Increasing structural strength and stiffness, for example, may reduce damage to displacement-sensitive components, but amplify floor accelerations, and in the process, increase

damage to acceleration-sensitive components. For both reasons, it is difficult to identify effective recovery-based design strategies and guide performance-objective selection.

A growing body of literature explores how optimization can be used to derive design insights using PBEE-based approaches (Rojas et al., 2019; Burton et al., 2019; Steneker et al., 2020; Steneker et al., 2022; Ghasemof et al., 2022; Karami et al., 2022; Amiri and Estekanchi, 2023). However, these studies typically focus on a single application (e.g., structural improvements to a building with a specific structural system), and do not demonstrate how such tools can be applied across different structural systems, building configurations, and upgrade strategies of interest. Reduced generality often stems from computational barriers associated with performing optimization, which requires tens to hundreds of thousands of building recovery performance realizations depending on the size of the design space. Lastly, none of these studies focus on controlling the functional recovery time of buildings, which is necessary to achieve community resilience goals.

To overcome these limitations, machine learning-based optimization methods can flexibly and rapidly isolate above-code building improvements to achieve recovery performance goals (Issa et al., 2023). This paper uses those methods to identify optimal improvement strategies across a large set of performance objectives. To understand how optimal strategies may change across different buildings, we repeat the analysis for two modern steel archetypes: one designed to code, and one in which the design drift limit is cut in half. The results illustrate the effect of reducing the drift limit on the necessary nonstructural improvement scope across all performance objectives explored, and the breadth of insights that are available when surrogate modelling is used in place of computational simulations.

2   Machine learning-based optimization framework

In this section, we describe the optimization framework used in the case study, which is adopted from Issa et al., (2023). The framework, which consists of five steps, can be used to isolate optimal recovery-based design improvements for a user-defined building, performance objective, and scope of improvements.

The framework (shown in Figure 1 for an intensity-based assessment) is split into two parts. Part I covers recovery modeling decisions along with design modification scoping. A baseline assessment is first performed to quantify the functional recovery time in absence of design improvements. Next, a performance objective, which defines the specific intensity and conditional reliability of interest, is defined (Equation 1).

$$P\big(T_{FR} > T_{target} \mid D_{FR}\big) < Y\% \qquad (1)$$

where: T_{FR} = functional recovery time
        T_{target} = target functional recovery time
        D_{FR} = selected hazard level for recovery-based design
        $Y\%$ = conditional probability of exceedance

Finally, building modifications (e.g., structural, nonstructural, utility backup) that will be used to achieve the performance objective are defined. The improvements considered in the design strategy are represented in an array of design parameters, X.

Part II covers the optimization procedure, which is responsible for isolating the optimal design solution X^* that minimizes a user-defined objective function. In this part of the framework, a feasibility check is first performed to ensure that the optimization algorithm will converge. If this is successful, then an optimization analysis using an objective function and algorithm of choice is performed to convergence. The generalized form of the optimization problem is shown in Equation 2.

$$\mathbf{X}^* = \arg\min\left(f(\mathbf{X})\right) \qquad (2)$$

with the constraints that 1) functional recovery time is less than the target, and 2) the design strategy $\mathbf{X}$ is in the design space.

Part I: Recovery modeling and scoping of design modifications
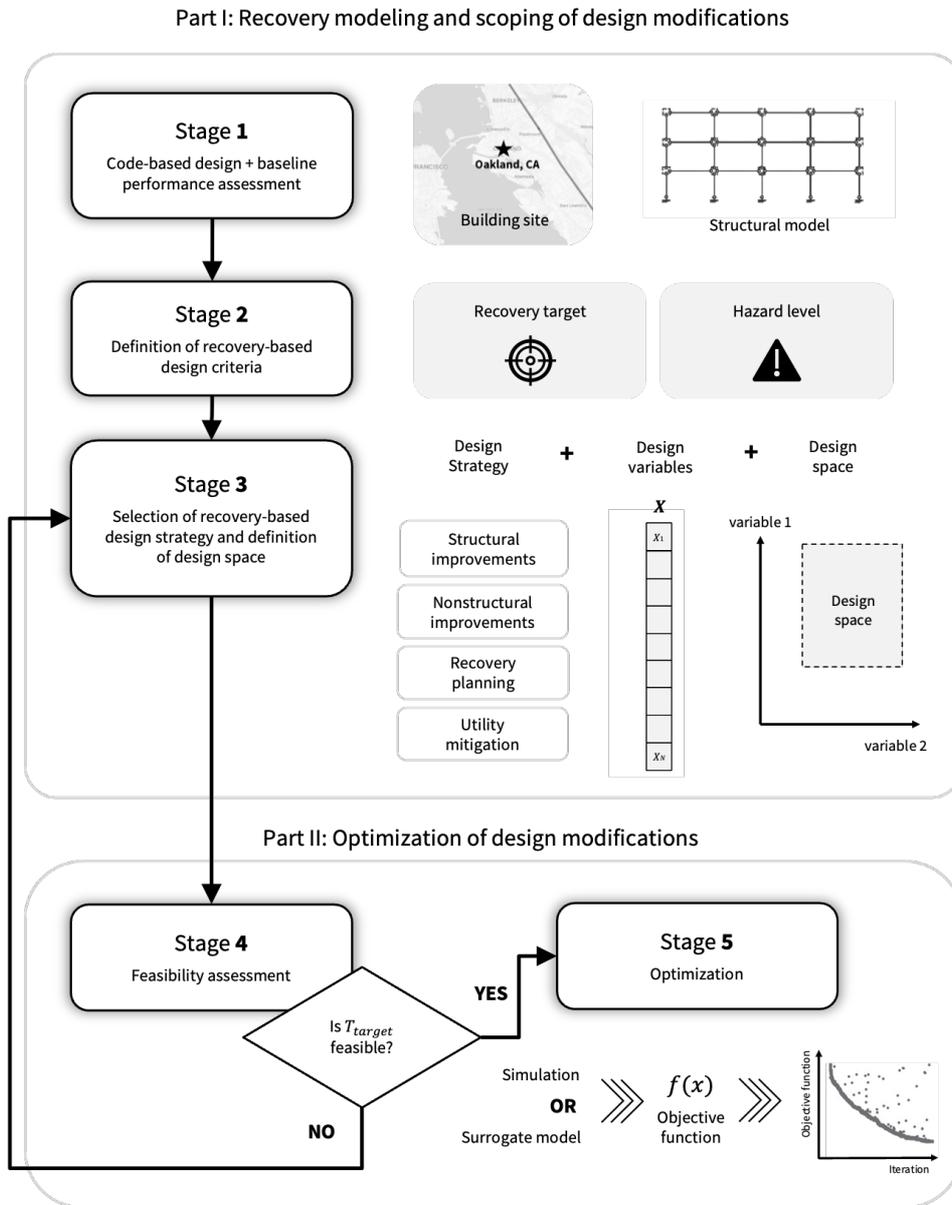

Figure 1: Overview of the recovery-based design optimization framework, applied to intensity-based performance objectives (Issa et al 2023).

Using surrogate models in place of direct FEMA P-58 and ATC-138 building recovery simulations (FEMA, 2018, Cook et al., 2022) during the optimization step (Stage 5 in Figure 1) can significantly reduce computational cost and enable exploration of the design space. As introduced in Issa et al., (2023), we use a surrogate model that de-couples structural and nonstructural improvements by including engineering demand parameters (EDPs) in the input feature vector. While this increases model complexity, it allows for use of the surrogate model across different hazard levels and structural upgrades without the need for retraining.

3 Case study using modern steel archetypes

3.1 Introduction

In this section, we apply the proposed framework to identify the optimal set of nonstructural component upgrades for two modern steel archetypes in Oakland, CA. The goal of this analysis is to study the implications of the selected performance objective on design. To facilitate this, we perform optimization analyses across a large set of performance objectives, spanning multiple target recovery times and hazard levels, repeated for each building. Hazard levels explored include the 72-, 224-, and 475-year return period.

3

Both building archetypes share the configuration shown in Figure 3 and are designed in accordance with ASCE 7-16 using the Equivalent Lateral Force (ELF) parameters in Table 1. The two buildings only differ in structural stiffness by varying the drift limit.
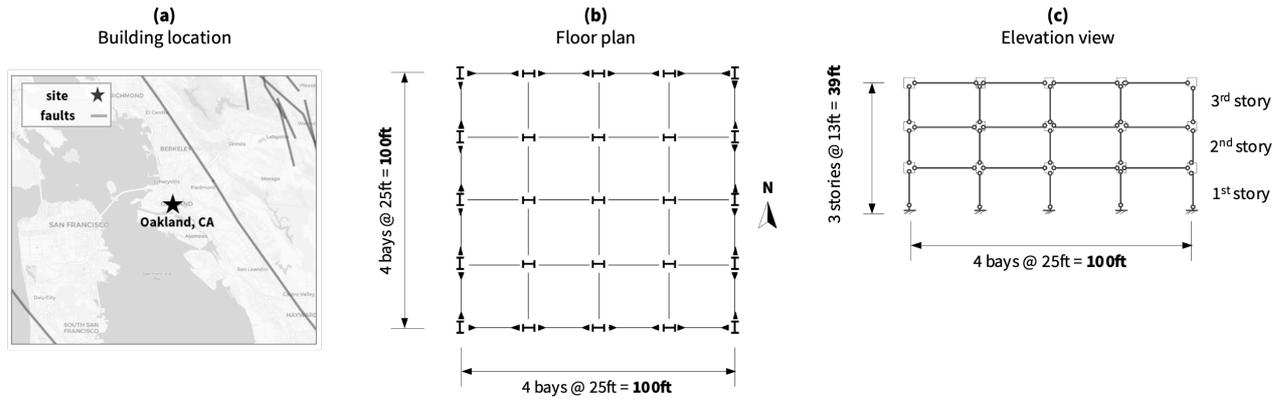
Figure 2: The location, floor plan, and elevation view of the three-story office building used in the case study.

Table 1: ASCE-7 ELF design parameters for each case study building.

	Site Class	Risk Category	Importance factor I_e	Drift limit	S_s	S_1	R	C_d
Archetype I	D	II	1.00	2.5%	1.559g	0.614g	8	5.5
Archetype II	D	II	1.00	1.25%	1.559g	0.614g	8	5.5

3.2 Application of optimization framework

Next, we describe the application of the framework shown in Figure 1, focusing on key decisions.

As part of Stage 1, each building is evaluated for baseline recovery performance (excludes any potential upgrades). We perform this assessment using a damage and loss assessment (FEMA P-58), followed by a downtime assessment (ATC-138) to compute functional recovery time.

For each archetype, we prepare an analytical structural model, a set of ground motions, and a building performance model. Analytical structural models are generated using Auto-SDA (Guan et al., 2020). Using the geometry and parameters in Figure 3 and Table 1, a OpenSees 2-D frame model for the X and Y direction is generated. For each building, 40 ground motions are selected using a Uniform Hazard Spectrum (UHS) target at the Oakland site, repeated for each return period considered in the case study. Performance models for office occupancy are directly adopted from the Issa et al., (2023) archetype, which consists of 41 components across nine building systems. The 50[th]-percentile baseline functional recovery times for each archetype are shown in Table 2. These times include any impeding factors that would prevent repairs from initiating.

Table 2: Baseline functional recovery times for each archetype, evaluated at three hazard levels of interest.

Hazard level	Archetype I (Drift limit = 2.5%)	Archetype II (Drift limit = 1.25%)
72-year	29 days	20 days
224-year	102 days	73 days
475-year	154 days	116 days

In Stage 2, the performance objective space is defined. To illustrate the breadth of performance objectives that can be explored using the framework, we consider target times spanning from 7 to 180 days for each of the hazard levels in Table 2. In each case, the conditional probability of exceedance is set to 50% as a proof of concept. While a lower probability of exceedance may be desired in real-world applications, using 50% allows us to compare the median functional recovery time directly with the target time.

As part of Stage 3, we explore nonstructural component improvements as the sole enhancement strategy for each archetype. Strengthening for each component (i.e., achieving a higher median capacity across all damage states) is modeled using the scheme introduced in Issa et al., (2023). Improvements to a single component are modeled by a scalar strengthening factor, x_i, which is applied to the fragility medians for all damage states. All original capacities for each component are derived directly from the FEMA P-58 fragility database. We assume that x_i varies continuously from $x_i = 1.00$ (no change) to $x_i = 3.00$ (triple the median capacity across all damage states). All values of x_i (one for each component) are organized in the array, $\mathbf{X_{NS}}$.

$$\mathbf{X_{NS}} = \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_N \end{bmatrix} \tag{3}$$

where N is the number of nonstructural components in the performance model. Component fragility functions are modified directly during each optimization evaluation using the values in $\mathbf{X_{NS}}$, and are used directly in recovery performance evaluations.

In Stage 4, we perform a feasibility assessment for each archetype and hazard level of interest to ensure that all optimization assessments converge for each performance objective (i.e., a viable design solution exists in each case). This is done by assessing the 50th percentile functional recovery time under maximal enhancement (in this case, tripling the median capacity of each component) for each archetype, and checking that the result less than 7 days, which is the minimum target time explored as part of the case study. We find that for each archetype, and each hazard level, an acceptable solution exists.

Finally, we perform the optimization as part of Stage 5 for each building. We select a cost-agnostic objective function $f(\mathbf{X})$ that measures the average nonstructural component improvement across all components in the performance model, which we assume to be an appropriate proxy for scope of work across the building (Equation 4). As part of the optimization, our goal is to minimize this, constrained by having the 50th percentile functional recovery time $T_{FR,50}$ meet the target time of interest.

$$f(\mathbf{X}) = \frac{1}{N} \sum_{x_i \in X_{NS}} x_i \tag{4}$$

s.t.

$$T_{FR,50} < T_{target}$$
$$x_i \in [1.00, 3.00] \; \forall \; x_i \in \mathbf{X_{NS}}$$

To reduce the computational cost associated with each evaluation of $f(\mathbf{X})$, we estimate $T_{FR,50}$ using a calibrated surrogate model that maps EDPs and building improvements $\mathbf{X_{NS}}$ directly to $T_{FR,50}$. We employ real-valued genetic algorithms for all analyses, since they work well in high-dimensional space, and are also compatible with simulation-based approaches (e.g., FEMA P-58 and ATC-138), which we use to validate select results.

3.3 Results

In this section, we discuss the results of the optimization for each archetype building, using the setup defined in Sections 3.1 and 3.2. Optimal designs in this section correspond to functional recovery target times from 7 to 180 days at the Oakland site with a conditional probability of exceedance of 50%. Optimal designs are generated for these objectives at the 72-, 224-, and 475-year hazard levels.

Optimal objective function values $f(X^*)$ for each building across the performance objective space are shown in Figure 4a. Here, $f(X^*)$ can be interpreted as a proxy for the nonstructural scope of work necessary to achieve the target time. Solid and dashed lines correspond to $f(X^*)$ values for Archetype I (at-code) and Archetype II (structurally improved), respectively.

To quantify the effect that the reduced drift limit has on scope, we estimate the value of $f(X^*)_{\Delta_{limit}=1.25\%} - f(X^*)_{\Delta_{limit}=2.5\%}$ in Figure 4b. When this difference is negative, structural improvements (i.e., cutting the drift limit in half) lead to a reduced nonstructural scope of work. When this difference is positive, the opposite is true. From this plot, we can quickly interpret performance objective ranges where reducing the drift limit has a positive and negative impact on the optimal nonstructural scope of work.
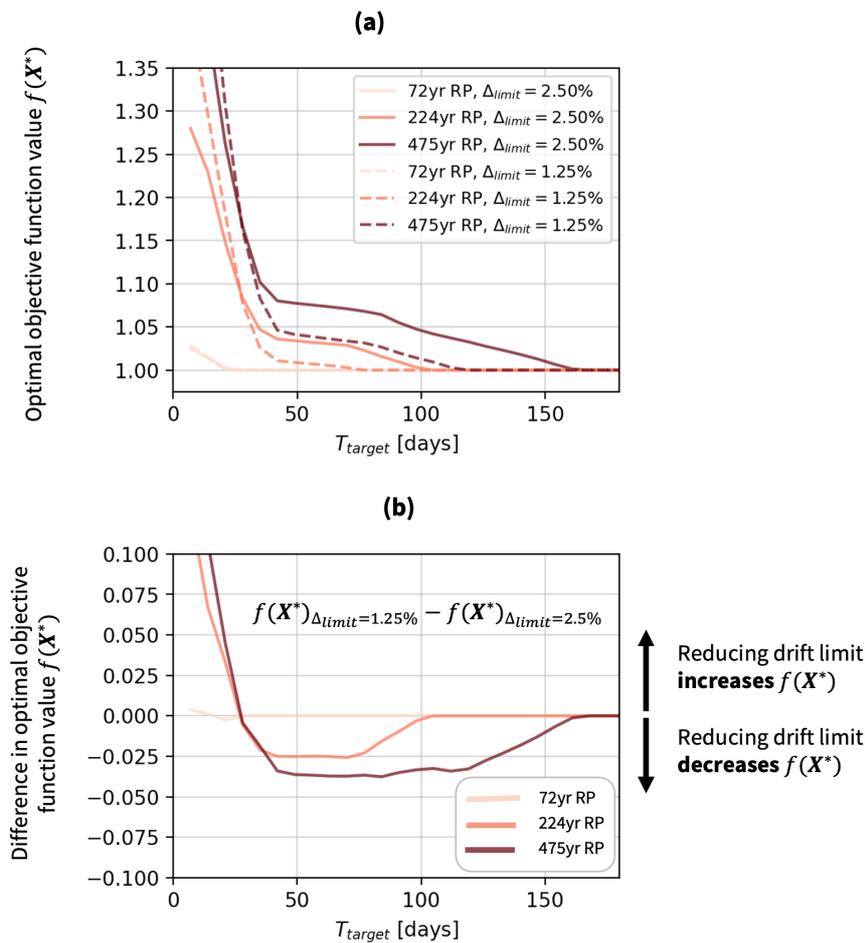

Figure 3: Optimal objective function (average NS hardening) values for each building, across different performance objectives (a) and the difference between objective function values for the two buildings (b)

A few observations can be made from the results in Figure 4a. First, we observe that for the majority of the target times and hazard levels explored, reducing the drift limit leads to reductions in $f(X^*)$. In general, these reductions are higher at higher hazard levels, which is also observed in Figure 4b. At the lowest intensity explored, we observe that the effect of reducing the drift limit is less pronounced and $f(X^*)$ values are roughly identical for both archetypes.

At higher hazard levels, we also observe that there are relatively large target time ranges in which reducing the drift limit warrants no additional nonstructural upgrades at all (i.e., $f(X^*) = 1.00$). For example, we observe that at the 475-year return period, $f(X^*) = 1.00$ from $T_{target} = 120$ days to $T_{target} = 155$ days when the drift limit is reduced. At the 224-year return period, the same is true from $T_{target} = 70$ days to $T_{target} = 120$ days.

Finally, we can isolate target times in which reducing the drift limit leads to an increase in $f(X^*)$, and hence, an increase in nonstructural scope. These regions are typically at very short target recovery times. At the 224 and 475-year return periods, this occurs when the target recovery time is less than 21 days. At the 72-year return period, this occurs when the target time is less than 14 days. An increase in $f(X^*)$ in these target time regions may be attributed to increased peak floor accelerations when the building is stiffened, which may increase the number of nonstructural components that require hardening.

These findings can also be deduced from Figure 4b, which illustrates the specific target time regions in which stiffening the building has a positive, negative, or negligible effect on optimal objective function values. The magnitude of the difference, along with the subset of target times in which stiffening is beneficial, both generally grow with increasingly large hazard levels. As observed in Figure 4a, at very low target times, reducing the drift limit triggers increasingly large objective function values. Such results suggest that at these target times, reducing the drift limit can be detrimental.

We draw several conclusions from these results. First, the breadth of the results illustrated in Figure 4 are made possible by surrogate modelling approaches that replace computationally expensive simulations during the optimization step. All analyses conducted in this case study were performed using a single surrogate model, since the input feature vector can accommodate different EDPs attributed to differing structural configurations and hazard levels of interest. Repeating this analysis using simulation-based approaches would prove to be computationally challenging in practice.

Second, the results clearly convey the impact of a structural improvement on the nonstructural scope of work for each hazard level. Due to the nonlinear relationship between building improvements and recovery performance, the degree to which structural improvements reduce nonstructural scope varies depending on both the hazard level and target time selected. Applying the framework to this problem helps to unveil trends that would otherwise be difficult to uncover using trial-and-error approaches.

Finally, the objective function values (particularly those in Figure 4b) can be used to guide performance objective selection, since they implicitly show how much nonstructural hardening is needed on average to achieve a given objective. These insights are prerequisites to supporting benefit-cost analysis studies, which quantify the benefits of reduced recovery time across the lifetime of the building to the costs associated with building upgrades. While the objective function in this case study is cost-agnostic, and the upper and lower bound constraints are likely not realistic for each component, refinements to both can make rapid and scalable benefit-cost studies a possibility when considering a broad set of performance objectives. Such studies can also be used to gauge the efficacy of specific recovery-based design strategies for the purposes of individual design and the development of provisions.

4   Conclusions

Identifying efficient recovery-based design strategies can be challenging in practice due to the computational cost of building recovery simulations, and the nonlinearity between these strategies and recovery performance. In this paper, we summarized our recent work to approach this problem using machine learning-based optimization and applied it to a new case study involving modern steel archetypes.

The framework was used to support the selection of recovery-based performance objectives by quantifying changes in optimal design strategies. For two building archetypes that differ in structural stiffness, we quantified the changes in the optimal objective function value: the average nonstructural hardening across an inventory of building components. This was repeated for three different hazard levels under consideration, and a set of target recovery times of interest.

Results illustrated specific target performance objectives where stiffening the building led to increased, decreased, and negligible changes in additional nonstructural upgrades. We found that at very low target recovery times, particularly at high hazard levels, the utility of reducing the drift limit is diminished, and can even be detrimental.

Finally, surrogate modeling, which maps building improvements directly to functional recovery time, is the enabling contribution of the framework that makes such broad analysis possible. These models significantly reduce the computational time associated with building recovery modeling and unlock the ability to perform optimization analyses at scale. Repeating this across a larger set of archetypes, building locations, and design strategies will improve recovery-based design intuition and support the development of future recovery-based provisions.

5   References

Ahmadie Amiri, H., Estekanchi, H. E. (2023). Life cycle cost-based optimization framework for seismic design and target safety quantification of dual steel buildings with buckling-restrained braces. *Earthquake Engineering & Structural Dynamics*.

Cook, D. T., Liel, A. B., Haselton, C. B., & Koliou, M. (2022). A framework for operationalizing the assessment of post-earthquake functional recovery of buildings. *Earthquake Spectra*, 38(3), 1972–2007.

Earthquake Engineering Research Institute (EERI) (2019). *Functional Recovery: A Conceptual Framework with Policy Options*.

FEMA. (2018). FEMA P-58: Seismic Performance Assessment of Buildings. Vol. 2– Implementation Guide. 2nd ed.

FEMA-USGS. (2023). FEMA P-366: Hazus Estimated Annualized earthquake losses for the United States.

Forell/Elsesser Engineers, Inc. (2018). UCSF Wayne and Gladys Valley Center for Vision | Forell/Elsesser Engineers, Inc.

Ghasemof, A., Mirtaheri, M., & Karami Mohammadi, R. (2022). Multi-objective optimization for probabilistic performance-based design of buildings using FEMA P-58 methodology. *Engineering Structures*, 254, 113856.

Guan, X., Burton, H., & Sabol, T. (2020). Python-based computational platform to automate seismic design, nonlinear structural model construction and analysis of steel moment resisting frames. *Engineering Structures*, 224, 111199.

HB-Risk. (2023). Seismic Performance Prediction Program (SP3).

Issa, O., Silva-Lopez, R., Baker, J. W., & Burton, H. V. (2023). Machine-learning-based optimization framework to support recovery-based design. *Earthquake Engineering & Structural Dynamics*.

Karami, M., Estekanchi, H. E., Hajirasouliha, I., & Mirfarhadi, S. A. (2022). Value-Based Seismic Performance Optimization of Steel Frames Equipped with Viscous Dampers. *Journal of Earthquake Engineering*, 1–27.

Mar D, Aher S. (2022). Affordable Resilience Using Rocking Walls and Foundation Dampers. *Proceedings of the 12th National Conference in Earthquake Engineering*.

Rojas, H. A., Foley, C., & Pezeshk, S. (2011). Risk-Based Seismic Design for Optimal Structural and Nonstructural System Performance. *Earthquake Spectra*, 27(3), 857–880.

SPUR. (2009). The resilient city: defining what San Francisco needs from its seismic mitigation policies. Technical report, San Francisco Planning + Urban Research Association.

Steneker, P., Filiatrault, A., Wiebe, L., & Konstantinidis, D. (2020). Integrated Structural–Nonstructural Performance-Based Seismic Design and Retrofit Optimization of Buildings. *Journal of Structural Engineering*, 146(8), 04020141.

Steneker, P., Wiebe, L., Filiatrault, A., & Konstantinidis, D. (2022). A framework for the rapid assessment of seismic upgrade viability using performance-based earthquake engineering. *Earthquake Spectra*, 38(3):1761-1787.

Zimmerman RB, Herdrich B. (2022). Resilience-based design of a seismically isolated, continued operations building. *Proceedings of the 12th National Conference in Earthquake Engineering*.

Zsarnóczay A, Deierlein G. (2020). PELICUN – A Computational Framework for Estimating Damage, Loss, and Community Resilience. *Proceedings of the 17th World Conference on Earthquake Engineering*.