14th International Conference on Structural Safety and Reliability - ICOSSAR'25
1-6 June, 2025 - Los Angeles California, USA
R. Ghanem, E. Johnson, S. Masri and A. Olivier (Eds)

SURROGATE-BASED SEISMIC RISK ASSESSMENT LEVERAGING SUPERVISED DATA-DRIVEN FEATURE EXTRACTION FOR THE EXCITATION CHARACTERIZATION

PARISA TOOFANI MOVAGHAR[1] AND ALEXANDROS A. TAFLANIDIS[2]

[1]University of Notre Dame
156 Fitzpatrick Hall, Notre Dame, IN 46556
ptoofani@nd.edu

[2]University of Notre Dame
156 Fitzpatrick Hall, Notre Dame, IN 46556
a.taflanidis@nd.edu

Key words: Seismic risk assessment; Surrogate modeling; Dimensionality reductions; Supervised feature extraction.

Abstract. Nonlinear response history analysis (NLRHA) has been established as the recommended approach for the vulnerability assessment of structural systems under seismic loading. To reduce the computational burden of this assessment, surrogate modeling techniques have emerged as a popular tool. Once trained based on a database of NLRHAs, the surrogate models (metamodels) can replace the original simulation model within the uncertainty propagation step (for example, using standard Monte Carlo analysis) to provide the desired vulnerability estimates. Unfortunately, the earthquake-acceleration time-series excitation creates a high-dimensional input that represents a significant challenge for any surrogate modeling approach, increasing the amount of data (NLRHAs in this context) that are needed for establishing high-accuracy metamodel-based predictions. Recent advancements suggest that dimensionality reduction (DR) techniques can enhance the efficiency of surrogate model estimates in such settings, though the accuracy of the formulation depends critically on the efficiency of the representation of the high-dimensional data in the established latent space. This study introduces a framework employing supervised feature extraction for DR, providing a compressed nonlinear latent representation of input features while retaining critical information relevant to structural responses. The input corresponds to both structural properties and the time-series earthquake excitation, while the output corresponds to the peak engineering demand parameters. The latent space representation of the input is subsequently utilized for efficient surrogate model training, enabling high-accuracy predictions of the targeted output even with limited number of simulations. The framework's effectiveness is validated using a dataset generated from a broad spectrum of ground motions, covering multiple intensities and a wide range of frequency content paired with diverse structural configurations. Results confirm that the proposed framework offers a significant reduction in computational demands and achieves reliable accuracy in probabilistic seismic reliability/vulnerability assessment.

1 INTRODUCTION

Accurate seismic risk assessment is critical for resilience planning at the individual building or regional scales. A key component of this assessments is the structural response prediction, typically expressed through the peak engineering demand parameters (EDPs) of interest under a wide range of excitation intensities[1]. For a comprehensive description of risk, these predictions need to be established using nonlinear response history analysis (NLRHA), providing samples of component-level responses that can ultimately enable high-fidelity damage, loss, and recovery quantifications[2]. At the same time the risk assessment needs to consider various types of uncertainties for the seismic hazard and the infrastructure models[1], to properly capture the potential variability in seismic demands. Despite computational advancements, addressing these uncertainties and achieving reliable seismic risk estimation poses a significant challenge[3], as it requires performing a large number computationally intensive NLRHAs. The challenge becomes even more pronounced at the regional scale, where the assessment of thousands of buildings significantly amplifies the computational burden[3].

One promising direction to address this challenge is the use of surrogate modeling techniques. Surrogate models, also referred to as metamodels or emulators, replace the original structural simulation model with a computationally efficient data-driven approximation, developed using a small number of judicially chosen simulations (i.e., training points). For seismic risk assessment applications, such surrogate modeling techniques can be used to emulate the peak EDPs (output) with significantly lower computational cost based on ground motion and structural parameter inputs. The key challenge for establishing surrogate models in this setting originates from the seismic excitation, as incorporating the full ground motion time histories as input to the surrogate model creates a high-dimensional problem, dramatically increasing the amount of data (NLRHAs in this context) that are needed for establishing high-accuracy predictive metamodels. One solution for addressing this challenge is to leverage stochastic emulators[4-7], treating the ground motion variability as aleatoric uncertainty source and describing the seismic excitation though a small number of intensity measures (IMs). Though efficient (require small only NLRHA databases for the metamodel development), this approach suffers from the fact that it can only emulate the EDP distribution under the aleatoric hazard uncertainty. It cannot establish predictions for specific ground motions, which might be necessary in a variety of settings, for example for capturing correlation across assets in regional risk assessment[3]. Another solution for addressing the high-dimensional excitation problem is to establish a physics-based dimensionality reduction (DR) of the earthquake ground motion time-series, and utilize a moderate number of pre-processed ground motion features[8-10], such as the response spectra, to describe each excitation. Such approaches require larger NLRHA databases for the metamodel calibration, and are typically coupled with surrogate models techniques that are appropriate for such databases, for example, deep neural networks (DNNs). Unfortunately, such DR approaches are not guaranteed to fully capture the temporal and frequency-dependent characteristics of ground motions that influence nonlinear structural response, reducing accuracy of the established surrogate models.

This paper addresses these gaps, and establishes a surrogate model for peak EDP predictions that explicitly utilizes ground motions as input and relies on a supervised data-driven feature extraction for DR for the excitation characterization. A variety of data-driven techniques[11], such as Principal Component Analysis (PCA), Kernel PCA (KPCA) and Autoencoders (AEs),

have been promoted over the years for DR of time-series data, established by extracting compact representations that retain the most informative patterns, rather than relying on the full raw signal. AEs have been shown[11] to be especially flexible and scalable, learning compact latent representations that capture both temporal and nonlinear dependencies and providing a practical balance between expressiveness and efficiency. Majority of relevant implementations have examined so far unsupervised DR, with sole objective to reduce the dimensionality of the time-series data. In contrast, the formulation examined here establishes a supervised feature extraction, considering explicitly the prediction of the peak EDPs. In addition, the integration of IMs as additional physics-based excitation features is considered in the DR.

2 PROBLEM FORMULATION FOR PEAK EDP APPROXIMATION

Consider structural model with associated vector of model parameters $\mathbf{x}_s \in \mathbb{R}^{n_s}$. Note that the proposed advancements are applicable to any type of infrastructure model (for example bridges), though to simplify the presentation the discussion will focus on building structures, which will be also the case considered in the illustrative implementation. Depending on the risk assessment application, i.e. whether this assessment pertains to (i) a specific asset or to (ii) a portfolio of assets within regional vulnerability estimation, this model may pertain to (i) an individual building or to (ii) a building archetype class. Vector $\mathbf{x}_s$ may pertain, respectively, to (i) the building properties that are considered as uncertain within the risk assessment (such as material properties) or to (ii) the parameters that distinguish the buildings within the portfolio (such as number of stories or fundamental period).

For assessing the seismic vulnerability of the building model through NLRHA, the seismic hazard is described through ground motion time-histories across different intensity levels. Let $\ddot{\mathbf{x}}_g \in \mathbb{R}^{n_g}$ denote the vector representing each discretized time-history input. Note that even though earthquake ground motions may have different durations, resulting in different lengths of their discretized representation, to support the data-driven surrogate model formulation a common length needs to be established for all (n_g dimensional vector in the notation adopted here). This can be seamlessly accommodated by padding zeros at the end of ground motion with shorter durations. Each of these ground motions can be characterized by an IM vector $\mathbf{x}_p \in \mathbb{R}^{n_p}$ representing the explanatory variables for the seismic hazard description[12], including quantities like the peak ground acceleration and velocity [PGA and PGV], or the spectral acceleration of a linear single degree of freedom (SDoF) oscillator with natural period T_s [$S_a(T_s)$]. Note that even though $\mathbf{x}_p$ may be equivalently interpreted[7] as a physics-based reduced-dimension description for $\ddot{\mathbf{x}}_g$, in the seismic hazard literature[12] it is predominantly considered as the vector or explanatory variables for the hazard. Different IMs may be considered for establishing $\mathbf{x}_p$, balancing between expressiveness, sufficiency and adequacy, with the appropriate choice dependent on the characteristics of the considered application[12, 13].

Finally, let $\mathbf{y} \in \mathbb{R}^{n_y}$ denote the EDP vector, corresponding to the peak responses required for the vulnerability quantification according to the damage/loss assessment module that is adopted for the performance quantification[14]. Objective of the surrogate modeling implementation considered here is to establish a data-driven approximation for $\mathbf{z}$ using as inputs $\mathbf{x}_s$ and $\ddot{\mathbf{x}}_g$, potentially supplemented through some choice for $\mathbf{x}_p$.

To accommodate the large dimension of $\ddot{\mathbf{x}}_g$ within this surrogate model formulation, a data-driven DR is considered, yielding mapping $\mathbf{z} = \boldsymbol{\varphi}(\ddot{\mathbf{x}}_g)$, with $\mathbf{z} \in \mathbb{R}^{n_z}$ representing the latent

ground motion features. The development of this mapping will be examined in detain in Section 3. The surrogate model for the peak EDP predictions is ultimately established as:

$$\hat{\mathbf{y}}(\mathbf{x}_s, \ddot{\mathbf{x}}_g) = \mathbf{f}(\mathbf{x}_s, \mathbf{z}, \mathbf{x}_p) \tag{1}$$

As shown in the notation used in Eq. (1), the approximation for $\hat{\mathbf{y}}$ utilizes inputs $\mathbf{x}_s$ and $\ddot{\mathbf{x}}_g$ but the surrogate model is established with respect to $\mathbf{x}_s$, $\mathbf{z}$ and potentially $\mathbf{x}_p$, with $\mathbf{z} = \boldsymbol{\varphi}(\ddot{\mathbf{x}}_g)$ and $\mathbf{x}_p$ derived based on $\ddot{\mathbf{x}}_g$.

The surrogate model approximation ultimately provides computationally efficient predictions of the EDP for any desired model configuration $\mathbf{x}_s$ and ground motion $\ddot{\mathbf{x}}_g$. This efficiency can be leveraged to support risk assessment using standard Monte Carlo (MC) techniques. For example, for seismic vulnerability assessment considering N_h seismicity scenarios with λ_o representing the occurrence rate of the *o*th scenario, the mean annual frequency (MAF) for the *j*th EDP for a specific building exceeding threshold **b** is calculated as:

$$\lambda(y_j > b \mid \mathbf{x}_s) = \sum_{o=1}^{N_h} \lambda_o P[y_j > b \mid \ddot{\mathbf{x}}_g^o, \mathbf{x}_s] \tag{2}$$

where $\ddot{\mathbf{x}}_g^o$ corresponds to the ground motion for the *o*th scenario, y_j to the *j*th EDP (*j* th component of vector **y**), and probability $P[y_j > b \mid \ddot{\mathbf{x}}_g^o, \mathbf{x}_s]$ can be readily estimated using the surrogate model approximation. If no additional sourced of uncertainty are included in the estimation of this probability then $P[y_j > b \mid \ddot{\mathbf{x}}_g^o, \mathbf{x}_s] = I[\hat{\mathbf{y}}(\ddot{\mathbf{x}}_g^k, \ddot{\mathbf{x}}_g) > \mathbf{b}]$ with *I*[] corresponding to the indicator function, taking value equal to 1 if the expression inside the brackets holds, else being zero. By leveraging the computational efficiency of the surrogate model a large number can be used for N_h in the hazard quantification (or in the MC simulations to address additional sources of uncertainty in $P[y_j > b \mid \ddot{\mathbf{x}}_g^k, \mathbf{x}_s]$) reducing statistical uncertainty while promoting a comprehensive risk quantification, leveraging NLRHA.

3 DEVELOPMENT OF SURROGATE MODEL FOR PEAK EDP PREDICTIONS

The formulation of the surrogate model for the peak EDP predictions is based on a training set with *n* observations obtained through NLRHA of the original structural model. Each training point in the set corresponds to a different ground motion and structural configuration pair. Superscript *i* will be used herein to describe the different training point and corresponding simulation results, with $\{\mathbf{x}_s^i, \ddot{\mathbf{x}}_g^i\}$ denoting the *i*th training point input pair and $\mathbf{y}^i$ the corresponding EDP output. Subscript *j* will be also used to denote the elements of a vector with y_j^i corresponding to the *j*th EDP for the *i*th training point. An appropriate IM vector is also considered for defining the explanatory excitation vector $\mathbf{x}_p$, with $\mathbf{x}_p^i$ denoting the corresponding value for the *i*th ground motion. The matrices with the simulation data (with different training points ordered along the columns) will be denoted with capital letters, with $\mathbf{X}_g \in \mathbb{R}^{n_g \times n}$ denoting the time-history data, $\mathbf{X}_s \in \mathbb{R}^{n_s \times n}$ denoting the structural configuration data, $\mathbf{X}_p \in \mathbb{R}^{n_p \times n}$ the explanatory variable data and $\mathbf{Y} \in \mathbb{R}^{n_y \times n}$ the EDP output data. The predictive model is composed of two components: (C.1) the feature extraction $\boldsymbol{\varphi}(\ddot{\mathbf{x}}_g)$ for mapping the high-dimensional ground motion time histories to the latent component vector $\mathbf{z}$ with dimension $n_z \ll n_g$; (C.2) the surrogate model $\mathbf{f}(\mathbf{x}_s, \mathbf{z}, \mathbf{x}_p)$ for establishing the final EDP approximation $\hat{\mathbf{y}}(\mathbf{x}_s, \ddot{\mathbf{x}}_g)$. The feature extractor can be established either on its own, with ultimate objective the approximation of $\ddot{\mathbf{x}}_g$, or it can be integrated within an end-to-end framework considering directly its impact on the accuracy of the $\hat{\mathbf{y}}(\mathbf{x}_s, \ddot{\mathbf{x}}_g)$ approximation.

3.1 Ground motion feature extractor

For the ground motion feature extractor some type of Neural Network (NN) architecture can be considered, widely shown[15, 16] to be efficient for DR related to time-series data. Candidate architectures include, auto-encoders (AEs), convolutional neural networks (CNNs), variational auto-encoders (VAEs), or recurrent neural networks (RNNs) like LSTMs. CNN-based approaches assume local translational invariance and emphasize temporal local trends[15], which may not align well with the non-stationary and globally distributed features of ground motion signals. LSTM-based architectures, while capable of capturing temporal dependencies, are often prone to overfitting on small datasets and require extensive training time, especially when the input sequence length is large[15, 17]—as is typical for seismic records. They also face significant challenges when the data do not share some strong underlying patterns[17], which will not hold when examining ground motion datasets across diverse seismic hazard intensities. Similarly, VAE frameworks impose probabilistic constraints on the latent space, which can introduce irrelevant variance[18]. Besides, they inherently suffer from a phenomena posterior collapse if the underlying distribution is not chosen wisely which makes their applicability for the application considered here challenging[18]. For these reasons AEs are promoted here for establishing mapping $\boldsymbol{\varphi}(\ddot{\mathbf{x}}_g)$.

AEs consists of two neural network components, an encoder $e(.)$ and decoder $d(.)$. The encoder maps the high-dimensional ground motion into the compact latent representation $\mathbf{z}$ [defining $\boldsymbol{\varphi}(\ddot{\mathbf{x}}_g)$] and the decoder reconstructs the signal from this latent space providing an approximation $\tilde{\ddot{\mathbf{x}}}_g$:

$$\tilde{\ddot{\mathbf{x}}}_g = d(e(\ddot{\mathbf{x}}_g)) \tag{3}$$

Traditionally, the AE objective is the reconstruction of the input, with training established using as loss function the reconstruction error ($\mathcal{L}_{rec}$) between $\ddot{\mathbf{x}}_g$ and its reconstruction $\tilde{\ddot{\mathbf{x}}}_g$:

$$\mathcal{L}_{rec} = \frac{1}{n \cdot n_g} \sum_{i=1}^{n} \sum_{j=1}^{n_g} \left(\ddot{x}_g^i(t_j) - \tilde{\ddot{x}}_g^i(t_j) \right)^2 \tag{4}$$

where $\ddot{x}_g^i(t_j)$ [or $\tilde{\ddot{x}}_g^i(t_j)$] refer to the time-series input [or its reconstruction] at time t_j for the *i*th earthquake, i.e. to the *j*th element of vector $\ddot{\mathbf{x}}_g^i$ [or $\tilde{\ddot{\mathbf{x}}}_g^i$].

However, minimizing reconstruction loss alone may lead to latent representations that retain irrelevant input features and/or overlook other critical features driving hysteretic structural behavior. For the application considered here, reconstructing a ground motion time history with high fidelity does not inherently have value. Small reconstruction errors may reflect accurate signal recovery, but do not guarantee that the extracted features improve the accuracy of the surrogate model for the EDP predictions. It is preferable to embed the DR within the surrogate model development, considering an end-to-end implementation. This is achieved by considering the loss function quantifying the surrogate model accuracy:

$$\mathcal{L}_{edp} = \frac{1}{n \sum_{j=1}^{n_y} \gamma_j} \sum_{i=1}^{n} \left(\gamma_j \sum_{j=1}^{n_y} \left(f_j\left(z^i, \mathbf{x}_s^i, \mathbf{x}_p^i\right) - y_j^i \right)^2 \right) \tag{5}$$

with γ_j representing the relative weight for the *j*th EDP. The overall loss function for the end-to-end implementation is established as a weighted combination of the reconstruction accuracy

and the EDP prediction performance:

$$\mathcal{L}_{total} = a\mathcal{L}_{rec} + (1-a)\mathcal{L}_{edp} \tag{6}$$

with weight a dictating the relative contribution from each of the component losses.

The following implementations can be distinguished based on the selection of a. (i) Unsupervised AE, denoted herein as UAE and corresponding to a=1. This formulation strictly focuses on reconstructing the ground motion signal. (ii) Supervised AE, denoted herein as SAE and corresponding to $0<a<1$. This formulation balances ground motion reconstruction and surrogate model response predictions. (iii) Physic-informed Latent Extractor, denoted herein as PLE and corresponding to a=1. This formulation ignores entirely the signal reconstruction, offering latent features optimized for the surrogate modeling. UAE is established independently of the surrogate model and therefore strictly utilizes $\mathbf{X}_g$ data, while SAE and PLE explicitly utilize the surrogate model response predictions, yielding a DR that is dependent on $\mathbf{X}_g$, $\mathbf{X}_s$ and $\mathbf{X}_p$ and is integrally related to the surrogate model development discussed next.

3.2 Surrogate model for the EDP predictions

The surrogate model $\mathbf{f}(\mathbf{x}_s, \mathbf{z}, \mathbf{x}_p)$ provides EDP predictions as a function of the structural $\mathbf{x}_s$ features, the latent ground motion representation $\mathbf{z}$ and the IM vector $\mathbf{x}_p$. To accommodate scalability to large dimensional inputs (and so large dimensional databases), while also promoting a seamless integration with the SAE and UAE DR implementations and NN-based formulation is examined utilizing a multi-layer perceptron[19] (MLP) feedforward NN, widely recognized to be capable of approximating complex nonlinear mappings between inputs and outputs. For input vector $\mathbf{x} = \{\mathbf{x}_s, \mathbf{z}, \mathbf{x}_p\}$ the MLP can be described as:

$$\mathbf{f}(\mathbf{x}) = W^{(k)}\sigma\left(W^{(k-1)}\sigma\left(...\sigma\left(W^{(1)}\mathbf{x} + b^{(1)}\right)...\right) + b^{(k-1)}\right) + b^{k} \tag{7}$$

where k is the number of layers, $W^{(k)} \in \Re^{d_k \times d_{k-1}}$ and $b^k \in \Re^{d_k}$ denote the learnable weight matrices and bias vectors at layer k, and σ is a nonlinear activation function such as ReLU (Rectified Linear Unit) or (Exponential Linear Unit) ELU[20]. This architecture is well-suited for capturing the intricate interactions between ground motion characteristics and structural configuration that govern seismic response behavior across different hazard intensities. Regularization techniques[19] such as dropout or L2 weight decay are included in the calibration to enhance generalization, particularly when training across ground motions representing diverse seismic intensities (and therefore degree of nonlinear behavior).

The MLP is trained by minimizing the supervised prediction loss $\mathcal{L}_{edp}$ in Eq. (5). When UAE is adopted as DR technique, the MLP calibration is established independent from the DR using strictly the minimization of $\mathcal{L}_{edp}$ as objective utilizing datasets $\mathbf{Y}$, $\mathbf{X}_s$, $\mathbf{X}_p$ and $\mathbf{Z}$, with the latter dataset obtained through the separate UAE implementation. When SAE or PLE is adopted as DR technique, then the MLP calibration is established in conjunction with the DR utilizing datasets $\mathbf{Y}$, $\mathbf{X}_s$, $\mathbf{X}_p$ and $\mathbf{X}_g$. Computationally the combined calibration is accomplished by an end-to-end manner utilizing the $\mathcal{L}_{total}$ loss function, where the parameters of both the DR component (where the input is $\mathbf{X}_g$) and the MLP regression head model (where the input is $\mathbf{X}_s$, $\mathbf{X}_p$ and $\mathbf{X}_g$) are jointly optimized through backpropagation using a unified loss function. This ensures that the DR prioritizes features that are most predictive of EDP.

4 ILLUSTRATIVE EXAMPLE

4.1 Structural and hazard model descriptions

The illustrative example considers a regional risk assessment implementation. The objective is the development of a surrogate model to estimate the peak drift response of five-story steel moment resisting frame building for the Western US. The multi-degree of freedom model (MDoF) for the archetype class follows the formulation presented in [21] and includes as structural characteristics (defining $\mathbf{x}_s$): the fundamental period T_1, the exponent c for defining the fundamental mode shape, dimensionless coefficients β_{if} and β_{ss} dictating the influence of infill and soft-stories, respectively, on the stiffness distribution along the height of the structure, the inter-story yield drift δ_y, the over strength ratio γ (describing an increase of peak strength relative to yield), the ductility ratio for the displacement at peak strength d, and the damping ratio ζ. This leads to n_s=8 dimensional input vector. Additional details for these parameter definitions are included in [21]. N_s=2,000 different structural configurations are established for $\mathbf{x}_s$, considering Latin Hypercube Sampling (LHS) in the domain of admissible values for these parameters considering the regional building portfolio. For the seismic hazard, ground motions are selected based on the PEER NGA-West2 database. A total of N_g=3,400 ground motions are considered, covering a wide range of intensity measures, durations, frequency content, and site conditions. Excitations are discretized using a time-step of 0.01s and a duration of 50 s leading to n_g=5,000.

4.2 Computational implementation details

The proposed framework consists of three candidate components: (1) the sequence-based feature encoder defining ultimately $\mathbf{z}$; (2) the decoder for input reconstruction, only needed if UAE or SAE are chosen as DR; and (3) the surrogate-modeling regression head for predicting EDPs. The feature encoder takes as input the ground motion time-series, and initially passes it through a fully connected (FC) layer of 1024 neurons with ReLU activation. A second FC layer with 512 neurons with ReLU activation follows, and the latent representation $\mathbf{z}$ is finally extracted via a dense layer with ELU activation. The latent dimension (n_z) is tunable. This is achieved iteratively, increasing its value until the relative improvement (in the loss function) plateaus. The decoder, when needed, reconstructs the original time-series input from this latent representation using two dense layers (512 and 1,024 neurons, respectively) followed by a final linear output layer. The latent vector $\mathbf{z}$ is concatenated with the vector of structural features $\mathbf{x}_s$ and, when desired, hazard IMs $\mathbf{x}_p$, and it is passed through two FC layers with ELU activations (with dimensions of latent size and half-latent size). The final regression layer outputs predicted EDPs (here peak inter-story drifts) using a linear activation function. All dense layers in the regression branch are regularized using L2 weight decay (with penalty 10^{-4}) to prevent overfitting and enhance generalization. The model is trained using the loss function $\mathcal{L}_{total}$ that might combine (for SAE) both $\mathcal{L}_{edp}$ and $\mathcal{L}_{rec}$. In the latter case, rather than preselecting weight a it is adaptively adjusted to represent the relative contribution of each loss component. The end-to-end model is optimized using the AdamW optimizer[22] with an initial learning rate of 0.001. An adaptive learning rate based on an exponential decay learning rate schedule[23] is used, where the learning rate is reduced by a factor of 0.9 every 10,000 training steps. The decay occurs in discrete steps (staircase decay). Early stopping was utilized, with convergence defined

as validation loss not improving over 10 consecutive epochs. The batch size remained fixed at 32, and weights are initialized using Xavier uniform initialization, with biases set to zero for numerical stability. A manual hyperparameter search was conducted to optimize model performance. The number of neurons, activation functions, dropout rates, and latent dimensions were systematically explored, with details shown in Table 1. When UAE is adopted as DR, then the AE and MLP are independently calibrated following the aforementioned guidelines.

Table 1: Hyper parameters selection candidates (in bold the chosen values are shown)

Hyper parameter	Feature Extraction Head	Regression Head
Number of Layers	[**5**, 7, 9]	[**2**, 3]
Hidden dimension	[64, 127, 256, **512, 1024**, 2048]	[**8, 16**, 32, 64]
Weight Decay	[$10^{-2,}$ 10^{-3}, **10^{-4}**]	N/A
Dropout Rate	[0.1, **0.15**, 0.2]	[**$10^{-2,}$** 10^{-3}]

4.3 Variants examined and calibration/validation setup

Different variants of the surrogate modeling framework are examined, considering the three different DR choices (UAE, SAE with a=0.1, and PLE as discussed in Section 3.1) and three cases for the $\mathbf{x}_p$ selection. The latter cases, with abbreviated notation in parenthesis are: (NH) no $\mathbf{x}_p$; (SA) the spectral acceleration at the structure's fundamental period $S_a(T_1)$; (RS) the response spectrum $S_a(T_s)$ for 250 periods within period range [0-5]s. These cases result in n_p values of 0, 1, and 250, respectively. Combination of DR and $\mathbf{x}_p$ selection yields a totals of nine variants but due to space limitations only selective cases of practical interest will be discussed further. The calibration set is established by combining each ground motion in the database with 2 different, randomly selected structural configurations, after entirely removing 20% of the ground motions and 20% of the structural realizations to serve as testing set for the validation. This results in a training set of n=5440. The validation set consists of n_v=1360 samples. Note that the formulation of the validation set, entirely composed of out-of-sample configurations for both the excitation and the structure, prevents information leakage.

4.4 Validation results and discussion

As validation statistics, the correlation coefficient (CC_j) and first-order normalized Wasserstein distance (W_j) are used, established by obtaining the statistics across testing sample (index j) for the comparison of the EDP predictions $\hat{y}_j(\mathbf{x}_s^i, \ddot{\mathbf{x}}_g^i)$ to the actual response $y_j(\mathbf{x}_s^i, \ddot{\mathbf{x}}_g^i)$. This is performed for each EDP with results reported per floor across EDP type (drifts or accelerations). Normalization for the Wasserstein distance is established using the standard deviation of the response $y_j(\mathbf{x}_s^i, \ddot{\mathbf{x}}_g^i)$. Figure 1 presents results for the three different DR variants for NH selection for $\mathbf{x}_p$, while Figure 2 presents results for the different $\mathbf{x}_p$ selection cases for the PLE variant. Results are only presented for three different floor EDPs due to space limitations. In Figure 2 another case is included, denoted as NGRS, for which the ground motion input in the DR is replaced with the full response spectrum (i.e., no features extracted based on the ground motion signal itself). Beyond the validation metrics presented in Figures 1 and 2, the latent feature size n_z is also of interest. For the cases examined in Figure 1 the n_z is estimated to be 256, 100 and 90 for UAE, SAE and PLE, respectively. For the cases examined in Figure 2 n_z decreases (for the PLE DR) from 90 (NH) to 64 for SA and 70 for RS.

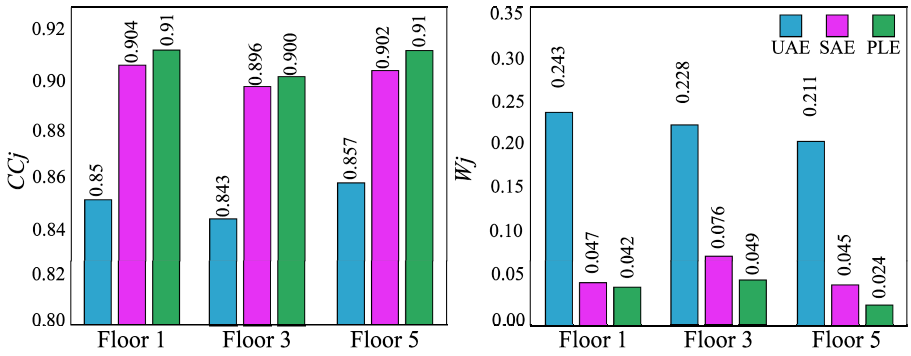

Figure 1: Validation comparison for different DR variants. Results (CC_j and W_j) presented for different floors.

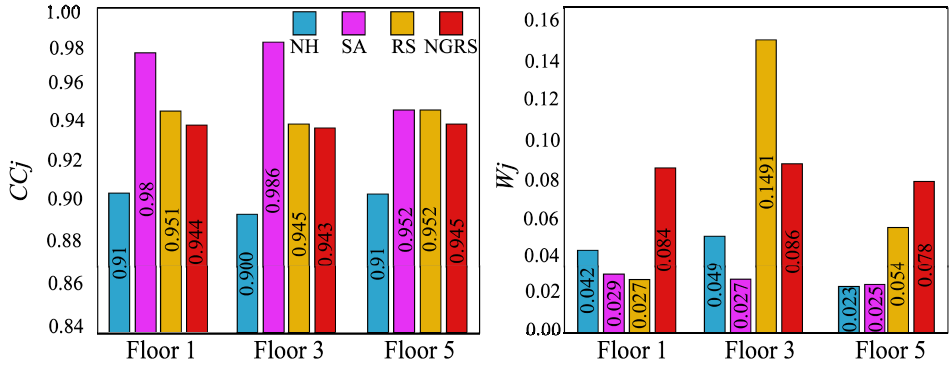

Figure 2: Validation comparison for different DR explanatory variable selections. Results (CC_j and W_j) presented for different floors.

Figure 1 clearly shows that unsupervised DR (UAE), focusing solely on the time-series earthquake signal reconstruction, is not appropriate for the implementation examined here, yielding substantially lower performance than SAE and PLE despite maintaining a larger number of latent feature (larger n_z as presented earlier). This validates our hypothesis that structural response information need to be explicitly considered within the DR. The best DR variant, yielding the best performance with respect to both CC_j and W_j across all floors, is PLE, showcasing that the signal reconstruction objective does not seem to provide any benefits within the DR. Ultimately, ground motions are inherently nonstationary and irregular, and the presence of repetitive or reconstructable patterns does not necessarily correlate with improvements in structural demand predictions. DR techniques prioritizing structural response predictability (i.e., PLE and SAE) are able to compress information more meaningfully (lower n_z value) than signal-reconstruction-based methods like UAE while supporting higher accuracy for the surrogate model.

Results in Figure 2 that a careful selection of additional explanatory variables $\mathbf{x}_p$ can clearly improve surrogate modeling performance, with SA outperforming NH by a considerable margins across both performance evaluation metrics and all EDPs while maintaining smaller number of ground motion features (smaller n_z value). The latter characteristic also clearly

shows that an appropriate $\mathbf{x}_p$ can infuse useful information in a response-centric DR, reducing the features that need to be retained to facilitate the predictive modeling. At the same time, including excessive information for $\mathbf{x}_p$ (RS case) may result in performance deterioration either due to overfitting (as also indicated by the larger n_z value compared to SA) or due to hindering the latent feature extraction by infusing a significant ground motion related information in the problem formulation to start with. Interestingly NGRS, which relies solely on reduced spectral features without direct ground motion information, shows improved performance over NH but fails to match the robustness of SA, highlighting that compressing the spectrum alone does not ensure retention of all meaningful structural response indicators. The results across both figures emphasize that a combination of a small number of informative IMs with ground motion based latent features extracted through a response-centric DR offers the greater accuracy and robustness for the peak EDP surrogate model predictions.

4.5 Vulnerability estimation demonstration

The validated surrogate model is now applied to perform seismic risk assessment for 5-story building structures in downtown LA. Risk is quantified through the MAF curves, estimated according to Eq. (2). Results for the actual MAF $\lambda_j(b) = \lambda(y_j > b \mid \mathbf{x}_s)$ are compared to the approximate MAF $\hat{\lambda}_j(b) = \lambda(\hat{y}_j > b \mid \mathbf{x}_s)$ established through the proposed surrogate modelling framework. Only the best performing variant is examined in this section for the latter. Results for the MAF curves for a specific structure are shown in Figure 3. Figure 4 presents results across a portfolio of 200 different buildings, focusing on the top story drift EDP response. Parts (a) and (b) of that figure shows the actual and approximated MAF curves for each of the buildings and the mean (in bold) across the buildings, while part (c) shows the histogram of the correlation coefficient of the MAF curves for each of the 200 buildings.

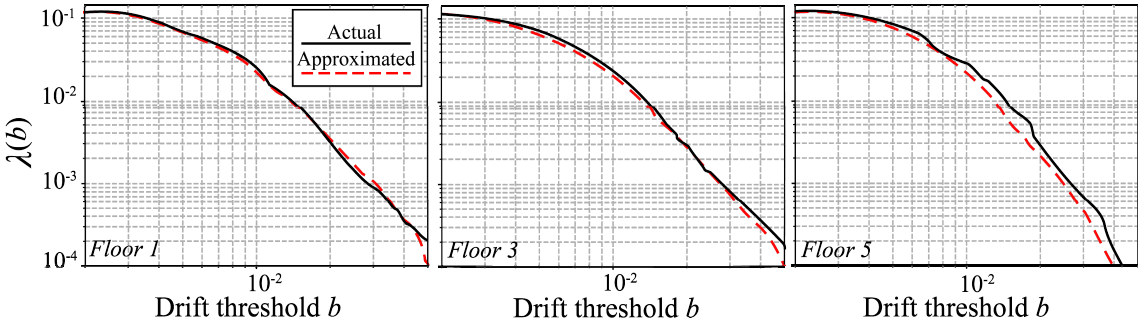

Figure 3: Drift EDP MAFs for a specific structure in downtown LA.

Results in Figure 3 show strong agreement between the predicted (red dashed lines) and actual (black solid lines) values, particularly in the mid-to-lower range of exceedance thresholds, which are most critical for performance-based seismic assessment. The consistency of the predicted MAF curves across multiple orders of magnitude—from frequent (10^{-1}) to rare (10^{-4}) events—demonstrates the model's capability in accurately capturing the probabilistic distribution of structural demand. Minor deviations are observed in the tail regions, which is expected given the data scarcity in these extremes. Results in Figure 4 show compatibility

between the MAF curves across the entire portfolio. This is evident by qualitatively comparing the swarms or the resultant average curve in parts (a) and (b) or by examining the correlation coefficient values in part (c). Overall, results across both figures confirm that the proposed surrogate framework preserves the underlying statistical behavior of structural response needed for risk quantification. Moreover, they demonstrate that (within the database validation (Section 4.3)) they can provide accurate information for the subsequent use of the surrogate model for vulnerability assessment, with similar degree of accuracy accomplished across both.

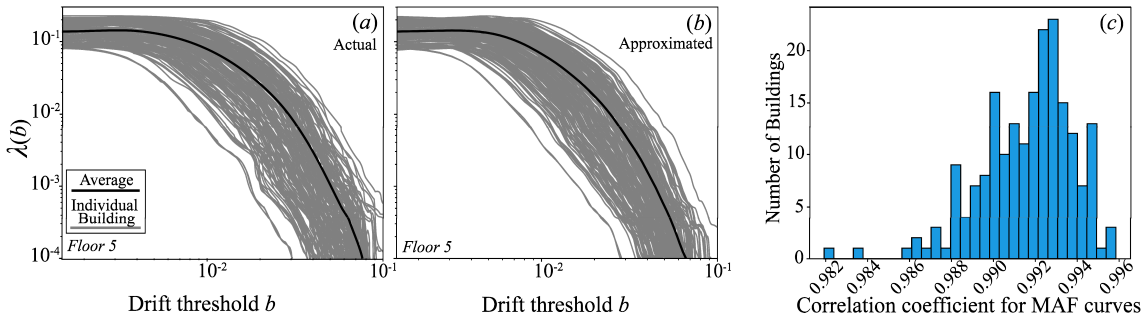

Figure 4: (a) Actual or (b) Approximate MAF curves for 200 different buildings in downtown LA along with the resultant mean MAF curve (across the buildings). (c) Histogram of the MAF curve correlation across each building.

5 CONCLUSIONS

This paper established a surrogate modeling framework for peak engineering demand parameter (EDP) predictions leveraging a data-driven feature extraction for the excitation characterization. The latent representation for the seismic input was derived in a supervised manner prioritizing its subsequent use to establish high-accuracy EDP predictions, instead of focusing on the input signal reconstruction. Additionally, the inclusion of excitation intensity measures (IMs) as inputs for the surrogate model was examined. The framework's effectiveness was validated utilizing a regional risk assessment setting. Results confirm that the supervised feature extraction and the judicious selection of IMs improve performance and that the calibrated surrogate model can subsequently support high-accuracy vulnerability predictions.

REFERENCES

[1] Baker JW, Cornell CA. Uncertainty propagation in probabilistic seismic loss estimation. Structural Safety. 2008;30:236-52.
[2] Bakalis K, Vamvatsikos D. Seismic fragility functions via nonlinear response history analysis. Journal of Structural Engineering. 2018;144:04018181.
[3] Deierlein GG, Zsarnóczay A. State of the art in computational simulation for natural hazards engineering. NHERI SimCenter. 2021.
[4] Kyprioti AP, Taflanidis AA. Kriging metamodeling for seismic response distribution estimation. Earthquake Engineering & Structural Dynamics. 2021;50:3550-76.
[5] Yi Sr, Taflanidis AA. Stochastic emulation with enhanced partial-and no-replication strategies for seismic response distribution estimation. Earthquake Engineering & Structural Dynamics. 2024;53:2354-81.

[6] Zhu X, Broccardo M, Sudret B. Seismic fragility analysis using stochastic polynomial chaos expansions. Probabilistic Engineering Mechanics. 2023;72:103413.
[7] Kim J, Wang Z. Uncertainty quantification for seismic response using dimensionality reduction-based stochastic simulator. Earthquake Engineering & Structural Dynamics. 2025;54:471-90.
[8] Kaffashchian MJ, Salkhordeh M, Karami Mohammadi R. A boosted deep learning-based approach for near real-time response estimation of structures under ground motion excitations. Structure and Infrastructure Engineering. 2024:1-20.
[9] Salkhordeh M, Alishahiha F, Mirtaheri M, Soroushian S. A rapid neural network-based demand estimation for generic buildings considering the effect of soft/weak story. Structure and Infrastructure Engineering. 2024;20:97-116.
[10] Wei B, Zheng X, Jiang L, Lai Z, Zhang R, Chen J et al. Seismic response prediction and fragility assessment of high-speed railway bridges using machine learning technology. Structures: Elsevier; 2024. p. 106845.
[11] Ashraf M, Anowar F, Setu JH, Chowdhury AI, Ahmed E, Islam A et al. A survey on dimensionality reduction techniques for time-series data. IEEE Access. 2023;11:42909-23.
[12] Baker J, Bradley B, Stafford P. Seismic hazard and risk analysis: Cambridge University Press; 2021.
[13] Kazantzi AK, Vamvatsikos D. Intensity measure selection for vulnerability studies of building classes. Earthquake Engineering & Structural Dynamics. 2015;44:2677-94.
[14] Goulet CA, Haselton CB, Mitrani-Reiser J, Beck JL, Deierlein G, Porter KA et al. Evaluation of the seismic performance of code-conforming reinforced-concrete frame building-From seismic hazard to collapse safety and economic losses. Earthquake Engineering and Structural Dynamics. 2007;36:1973-97.
[15] Casolaro A, Capone V, Iannuzzo G, Camastra F. Deep learning for time series forecasting: Advances and open problems. Information. 2023;14:598.
[16] Franceschi J-Y, Dieuleveut A, Jaggi M. Unsupervised scalable representation learning for multivariate time series. Advances in Neural Information Processing Systems. 2019;32.
[17] Bai S, Kolter JZ, Koltun V. An empirical evaluation of generic convolutional and recurrent networks for sequence modeling. arXiv preprint arXiv:180301271. 2018.
[18] He J, Spokoyny D, Neubig G, Berg-Kirkpatrick T. Lagging inference networks and posterior collapse in variational autoencoders. arXiv preprint arXiv:190105534. 2019.
[19] Kruse R, Mostaghim S, Borgelt C, Braune C, Steinbrecher M. Multi-layer perceptrons. Computational intelligence: a methodological introduction: Springer; 2022. p. 53-124.
[20] Clevert D-A, Unterthiner T, Hochreiter S. Fast and accurate deep network learning by exponential linear units (elus). arXiv preprint arXiv:151107289. 2015.
[21] Yi S-r, Taflanidis AA, Toofani Movaghar P, Galasso C. Impact of structural information fidelity on reduced-order model development for regional risk assessment. Structural Safety. 2025:doi.org/10.1016/j.strusafe.2025.102602.
[22] Loshchilov I, Hutter F. Decoupled weight decay regularization. arXiv preprint arXiv:171105101. 2017.
[23] Li Z, Arora S. An exponential learning rate schedule for deep learning. arXiv preprint arXiv:191007454. 2019.