

ORIGINAL ARTICLE

Comparison of surrogate strategies for damage-aware optimization of buckling-delayed shear-link dampers with adaptive finite element validation

J. Irazábal¹ | J. Ramírez^{1,2} | J. M. González^{1,2} | L. Lázaro¹ | F. Rastellini^{1,2} | G. Bozzo^{2,3} | L. Bozzo⁴

¹Centre Internacional de Metodes Numerics en Enginyeria (CIMNE), Barcelona, Spain | ²Universitat Politècnica de Catalunya (UPC), Barcelona, Spain | ³SLB Devices, Barcelona, Spain | ⁴Luis Bozzo Estructuras y Proyectos S.L., Barcelona, Spain

Correspondence: J. Irazábal. jirazabal@cimne.upc.edu

Received: Date Month Year | **Revised:** Date Month Year | **Accepted:** Date Month Year

Keywords: Buckling-delayed shear link, seismic energy dissipation, FEM validation, surrogate modelling, machine learning, radial basis functions, Differential Evolution.

Abstract

Buckling-delayed shear-link dampers are used in seismic-resistant structures as passive devices that concentrate energy dissipation while limiting damage to the primary system. Their geometric optimization requires balancing high energy dissipation with strict control of local damage. Finite element models can accurately reproduce the nonlinear cyclic response of these devices and provide internal quantities such as damage indicators and local distortion, but their computational cost prevents their direct use within iterative optimization loops. This work proposes an adaptive surrogate-assisted optimization framework for buckling-delayed shear-link dampers. First, experimentally calibrated nonlinear finite element models are used to generate reference datasets for dampers with different geometric and mechanical configurations. Supervised learning models are initially evaluated, with support vector regression and Gaussian process regression consistently providing high predictive accuracy. This suggests that kernel-based, distance-dependent approximations are well suited to the problem, motivating the introduction of radial basis function surrogates as a computationally efficient alternative. The surrogate predictions are coupled with a differential evolution algorithm through a damage-aware objective function that controls local damage and uses dissipated energy as a performance criterion. Optimized geometries are finally re-evaluated with finite element simulations. When the surrogate error exceeds the adopted tolerances, the new simulation result is added to the dataset and the surrogate models are retrained. The proposed framework provides an efficient damage-aware optimization of seismic energy dissipation devices.

1 | Introduction

Shear-link beam (SLB) dampers are passive energy dissipation devices widely used in seismic-resistant structures, designed to undergo stable inelastic deformations while limiting damage to the primary system. Their configuration concentrates inelastic demand in replaceable components, enabling high energy dissipation, ductile response and stable hysteretic behavior under severe cyclic loading^{1,2}. Within this family, buckling-delayed shear-link (BDSL) dampers incorporate a mechanical configuration that promotes shear-dominated behaviour while delaying local and global buckling.

The global hysteretic response of BDSL dampers can be characterized through experimental testing. However, laboratory campaigns present important limitations when addressing design optimization problems: internal state variables such as local plastic strains, stress triaxiality or damage indicators cannot be directly measured and the high cost and logistical complexity of experimental programs restrict the number of geometric configurations that can be explored. These limitations become even more critical when the design objective is not only to increase global energy dissipation and maximum displacement supported by the device, but also to control where and how damage develops within the device, making systematic optimization impractical.

In this context, finite element method (FEM) simulations provide a powerful framework for generating high-fidelity datasets across a wide range of geometric configurations, making them an ideal foundation for surrogate-based optimization strategies. Advanced nonlinear FEM models can accurately reproduce experimental cyclic behaviour, including plasticity, geometric nonlinearity, contact interactions, local instability and damage evolution, while providing access to both global response quantities and local indicators governing failure mechanisms. However, their high computational cost makes their direct use within optimization loops inefficient. This limitation motivates the use of surrogate models trained on FEM-generated data, which can approximate the structural response with reduced computational effort. The use of surrogate models enables the possibility of evaluating a large number of design configurations, facilitating systematic optimization.

Over the past decades, FEM has been widely used to study BDSL dampers and other passive seismic energy dissipation devices, providing detailed insight into nonlinear cyclic response, stiffness degradation and local inelastic mechanisms^{3,4}. It has also supported the optimization and parametric analysis of these devices by enabling systematic exploration of geometric configurations and performance criteria under prescribed loading^{5,6}. In this context, FEM-based studies have been extensively applied to characterize the mechanical response of metallic dampers. Motamedi et al.⁷ investigated accordion metallic dampers through combined experimental and numerical analyses, assessing the influence of key geometric variables on stiffness, strength and energy dissipation. Ghamari et al.⁸ studied I-shaped shear links in concentrically braced frames, while Xiong et al.⁹ examined replaceable steel shear links with different short-length ratios, highlighting the strong influence of geometry on cyclic performance and failure modes. Simplified analytical and semi-empirical models have also been proposed to reduce computational cost¹⁰ and more recent simulation-based studies have further explored the role of geometric and material variables in damper performance¹¹.

Geometric optimization has also been extensively explored. Zhang et al.¹² proposed a Kriging-assisted framework to maximize hysteretic energy in coupling beam dampers. Farzampour et al.¹³ optimized butterfly-shaped shear links by maximizing the ratio between dissipated energy and plastic strain, while Khatibinia et al.^{14,15} developed efficient strategies for U-shaped dampers using FEM and surrogate models. Shi et al.¹⁶ introduced a non-parametric shape optimization framework for shear panel dampers and Saleh et al.^{17,18} extended this line through topology optimization of shear-link configurations. More recent contributions include the hybrid cellular automata approach by Mendoza-Cuy et al.¹⁹ and the statistical optimization framework by Rios et al.²⁰.

Data-driven approaches have mainly focused on response or property prediction. Chan et al.²¹ used nonlinear autoregressive exogenous (NARX) models to reproduce hysteretic behaviour. Bae et al.²² developed models for low-cycle fatigue estimation and Almasabha et al.²³ predicted the shear strength of short steel links using ML. Elgammal et al.²⁴ modelled hysteretic restoring forces using data-driven approaches, while Hu et al.²⁵ proposed explainable ML models for the probabilistic prediction of buckling stress. Physics-informed approaches have also been explored, such as the PINN framework proposed by Hu et al.²⁶.

All these works demonstrate the increasing interest in applying FEM-based and data-driven approaches, as well as in combining both, to analyse, understand and optimize seismic energy dissipation devices. However, most of these studies focus either on the prediction of the hysteretic response or on maximizing energy dissipation, leaving a critical aspect insufficiently explored: the need to control local damage while maintaining adequate dissipative capacity. In practice, excessive local damage may compromise structural integrity, reduce durability and lead to premature failure, even when global energy dissipation is improved.

The present work addresses this gap through a damage-aware surrogate-assisted optimization methodology in which the objective is not only to maximize distortion or energy dissipation, but also to balance dissipative performance with damage indicators derived from FEM simulations. The proposed approach combines: (i) experimentally calibrated nonlinear FEM models used as numerical ground truth; (ii) supervised surrogate models trained to predict local damage and distortion indicators; (iii) a Differential Evolution (DE) optimizer; and (iv) an adaptive FEM validation and retraining loop.

Figure 1 summarizes the proposed workflow. The different stages of the methodology, together with the surrogate modelling, optimization strategy, validation procedure and corresponding results and conclusions, are described in the following sections.

2 | Buckling-delayed shear-link damper

The BDSL dampers analysed in this work, with one representative configuration shown in Figure 2, are designed to concentrate energy dissipation in localized reduced-thickness zones, hereafter referred to as windows, while preserving the structural integrity of the surrounding frame. The optimization variables correspond to the window thicknesses, whereas the frame dimensions are kept fixed. The dissipative element is connected to a load-transfer system through a mechanism that allows imposed in-plane displacement while preventing axial force transmission, thereby promoting a shear-dominated response. Under cyclic loading, plastic deformation is intended to concentrate in the windows, whereas the frame provides load transfer and stability. This separation of functions leads to a non-trivial design problem. Thin windows may enhance ductility and dissipative activation, but they may also promote excessive damage localization. Conversely, thicker windows may increase strength while transferring inelastic demand to the frame. Since severe frame damage may compromise the structural integrity of the device, frame damage must be penalized more strongly than window damage. At the same time, the dissipative demand should be distributed as uniformly as possible among the windows, avoiding configurations in which a single window absorbs most of the deformation while the remaining windows stay underused. Consequently, the design problem cannot be reduced to maximizing force or total dissipated energy alone, but must also control where damage develops and how the windows participate in the dissipative process.

The design variables considered in this work are the window thicknesses,

$$\mathbf{x} = [t_{w,1}, t_{w,2}, \dots, t_{w,W}], \quad (1)$$




FIGURE 1 | Flow chart of the proposed adaptive surrogate-assisted optimization framework.



FIGURE 2 | Representative BDSL damper configuration considered in the optimization.




FIGURE 3 | Geometric scheme of a representative BDSL device considered in this study.

where W denotes the number of windows. The device width and height are identified by B and H , respectively. Five geometry families are analysed, defined by different combinations of height, width and number of windows. Due to confidentiality constraints, the main dimensions are reported in normalized form with respect to reference values, H_{ref} and B_{ref} . Figure 3 shows the geometric scheme of one representative BDSL configuration. The example corresponds to a two-window device; throughout the paper, windows are numbered from top to bottom.

The admissible thickness ranges are defined for each geometry family according to manufacturing constraints. The main characteristics of the analysed families are summarized in Table 1. In all cases, the frame thickness is kept constant.

3 | Validation of the FEM numerical model

The surrogate models developed in this work were trained using data generated from three-dimensional FEM simulations. The numerical model is based on a previously calibrated and validated representation of the BDSL device, described in detail in Ramirez et al.²⁷. An example of the numerical setup is shown in Figure 4. The simulations were carried out using the COMPACK code, an explicit dynamic FEM solver based on an incremental formulation²⁸. The model accounts for large displacements, material and geometric nonlinearities, contact interactions and the boundary conditions associated with the experimental configuration.

The dissipative steel component is modelled as ASTM A36 steel, whose cyclic plastic behaviour is represented by the Yoshida-Uemori model^{29,30}. This constitutive law allows the model to reproduce cyclic hardening, softening and Bauschinger-type effects under large plastic deformation. The steel component is

TABLE 1 | Geometry families considered in the optimization and admissible window thickness ranges.

Family	H/H_{ref}	B/B_{ref}	W	Design variables	Thickness bounds
F_1	1.00	1.00	2	$t_{w,1}, t_{w,2}$	10–20 mm
F_2	1.00	1.17	2	$t_{w,1}, t_{w,2}$	10–20 mm
F_3	1.50	1.00	3	$t_{w,1}, t_{w,2}, t_{w,3}$	5–14 mm
F_4	1.50	1.17	3	$t_{w,1}, t_{w,2}, t_{w,3}$	5–14 mm
F_5	2.00	1.17	5	$t_{w,1}, \dots, t_{w,5}$	5–12 mm

discretized using linear eight-node hexahedral solid elements, providing a structured three-dimensional mesh suitable for extracting local stress, strain and damage-related fields.

The imposed displacement is applied through an actuator-like connector that transfers horizontal displacement while preventing axial load transmission, thereby reproducing the kinematic condition required to promote a shear-dominated response. Three experimental tests were performed following the loading protocols defined in the ANSI/AISC 341-16³¹ and ASCE 7-22³² qualification standards for seismic energy dissipation devices. Those protocols involve cyclic loading with progressively increasing amplitudes. Additional contact and confinement conditions are included to represent the experimental anti-buckling system and to control local and global instability.

The model was calibrated and validated against cyclic experimental tests performed on representative BDSL specimens. The calibration involved the material parameters, assembled geometry, contact definitions, support flexibility and boundary conditions. The validated model accurately reproduces the main global experimental responses. Figure 5 shows the comparison between experimental and numerical results, confirming the suitability of the FEM model as a numerical reference for configurations beyond those experimentally tested.

4 | Dataset generation and surrogate modelling

Once validated, the FEM model is used to generate the datasets required for surrogate training and optimization. The optimization strategy relies on local damage indicators in the dissipative windows and surrounding frame, together with local distortion measures associated with the activation of the dissipative mechanism. Since these quantities are difficult to measure experimentally, the use of a high-fidelity FEM model provides valuable access to the internal state variables and local fields governing damage evolution and energy dissipation.

The damage indicator adopted in this work is the Triaxial Failure Damage Map (TFDMap)³³. This stress-triaxiality-based indicator evaluates the proximity of each material point to ductile failure by comparing its stress triaxiality and accumulated equivalent plastic strain with a reference failure envelope^{34,35,36,37}. In this study, the TFDMap is used as a post-processing damage-screening indicator, not as a constitutive fracture criterion. Its purpose is therefore not to explicitly predict crack initiation, but to compare geometrical configurations and ensure that optimized designs remain within acceptable damage levels.

4.1 | Design of experiments

The FEM campaign is planned to cover the admissible parameter domain of each device family while ensuring a homogeneous exploration of the multidimensional space. The design variables are the window thicknesses $t_{w,i}$, whose combinations are generated through a Design of Experiments (DoE) strategy based on Latin Hypercube Sampling (LHS) optimized with the maximin criterion³⁸. This approach provides a near-random yet space-filling distribution of samples, reducing clustering and improving the representation of the admissible domain.

For each geometry family, the optimization cycle starts from an initial set of FEM simulations selected to provide reasonable coverage of the design space while keeping the computational cost as low as possible. The number of initial samples is defined according to the dimensionality of each family: 8 samples for two-window devices, 16 samples for three-window devices and 64 samples for five-window devices. In all cases, the number of samples is selected as a power of two, which facilitates the potential use of Progressive Latin Hypercube Sampling (PLHS)³⁹ in future iterations. This would allow the DoE to be expanded while preserving its space-filling properties and avoiding the need to repeat previously computed simulations.

To improve surrogate robustness near the admissible limits, the sampling domain is extended slightly beyond the actual optimization ranges reported in Table 1. This reduces the risk of extrapolation when evaluating candidate solutions close to the true design limits. The extended thickness ranges and the number of initial simulations for each family used in the DoE are summarized in Table 2.

For every sampled configuration, a FEM simulation is performed under a displacement-controlled cyclic loading protocol. Since the admissible deformation demand depends on the size of the device, different loading patterns are adopted according to the device height. As shown in Figure 6, the maximum displacement amplitude increase with the device height, consistently with the expected performance range of each geometry family.

The resulting dataset stores the input variables and the structural response quantities extracted at the final time of each simulation. The target outputs include damage indicators in the windows and in the frame, together with local distortion measures associated with dissipative activation. For compactness in the optimization formulation, the aggregated TFDMap indicator in window i is denoted by $\mathcal{D}_i$, whereas the corresponding frame indicator is denoted by $\mathcal{D}_f$. Both quantities are computed as the average TFDMap value of the 12 nodes with the highest values within each region. This allows to capture the most critical damage levels while reducing sensitivity to isolated numerical peaks.




FIGURE 4 | FEM validation model of the BDSL device: mesh discretization, main components, boundary conditions and local/global buckling control.




FIGURE 5 | Experimental-numerical validation of the BDSL model: hysteretic response (left) and cumulative dissipated energy (right).

The maximum local shear distortion in each window is denoted by $\varepsilon_{xy,i}$ and is used as an indicator of the energy dissipation capacity. When included in the objective function, each window contribution is weighted according to its effective geometric volume, so that the optimization accounts for both local strain intensity and the material volume involved in the dissipation process.

4.2 | Supervised ML surrogate models

This work compares several supervised surrogate models for predicting FEM-derived damage and distortion indicators. The considered models cover three families: tree-based methods, including Random Forest (RF)⁴⁰, Gradient Boosting Regression (GBR)⁴¹ and XGBoost⁴²; kernel-based methods, including Support Vector Regression (SVR)⁴³ and Gaussian Process Regression

(GPR)⁴⁴; and neural-network models, represented by the Multi-layer Perceptron (MLP)^{45,46}. RF relies on bootstrap aggregation of decision trees, GBR and XGBoost are sequential boosting approaches, SVR and GPR exploit kernel functions to model nonlinear relationships and MLP approximates nonlinear input-output mappings through interconnected layers.

For every geometry family, a separate regression model is trained for each target output. The input vector contains the window thicknesses of the corresponding device, while the outputs are the local distortion indicators in each window, $\varepsilon_{xy,i}$, the window damage indicators, $\mathcal{D}_i$, and the frame damage indicator, $\mathcal{D}_f$. Therefore, for a geometry family with W windows, $2W + 1$ surrogate models are trained: W models for window distortion, W models for window damage and one model for frame damage.

TABLE 2 | Geometry families, window thickness ranges and number of simulations considered during the DoE generation.

Family	H/H_{ref}	B/B_{ref}	W	Design variables	Thickness bounds	Simulations (DoE)
F_1	1.00	1.00	2	$t_{w,1}, t_{w,2}$	8–22 mm	8
F_2	1.00	1.17	2	$t_{w,1}, t_{w,2}$	8–22 mm	8
F_3	1.50	1.00	3	$t_{w,1}, t_{w,2}, t_{w,3}$	4–16 mm	16
F_4	1.50	1.17	3	$t_{w,1}, t_{w,2}, t_{w,3}$	4–16 mm	16
F_5	2.00	1.17	5	$t_{w,1}, \dots, t_{w,5}$	4–14 mm	64




FIGURE 6 | Displacement-controlled cyclic loading patterns adopted for the different device heights considered in the FEM campaign.

For each output, all candidate algorithms are evaluated independently and the selected model is retained for the optimization stage.

Hyperparameters are optimized using Bayesian optimization⁴⁷, with 40 evaluations per model and RMSE as the refit criterion. The first 10 evaluations are randomly sampled to explore the search space, while the remaining 30 are guided by the Bayesian surrogate model. This strategy provides a more efficient alternative to exhaustive grid search, particularly considering the number of geometry families, target outputs and candidate algorithms analysed. It also allows a broader exploration of continuous hyperparameter ranges. The hyperparameter search spaces considered are summarized in Appendix A.

The cross-validation strategy is adapted to the dataset size. Leave-One-Out validation is used for $N \leq 20$, repeated four-fold cross-validation with five repetitions for $21 \leq N \leq 80$, and shuffled five-fold cross-validation for larger datasets. For small datasets, the search spaces of tree-based models are additionally restricted to reduce overfitting.

Model selection is performed in two stages. First, for each candidate algorithm, Bayesian hyperparameter optimization is carried out using cross-validated Root Mean Squared Error (RMSE) as the refit criterion. The best hyperparameter configuration is therefore the one with the lowest mean RMSE. Then, the best configurations from all candidate algorithms are compared. The model with the lowest mean RMSE defines a competitive threshold and all models with an RMSE within 5% of this value are retained. Among these competitive models, the final selection is based on the lowest relative RMSE dispersion, computed as the standard deviation of the fold-wise RMSE divided by the mean RMSE. If two models have the same dispersion, the one with the lower mean RMSE is preferred. This procedure, summarized in Figure 7, favours surrogates that are both accurate and stable.

Preliminary executions of the proposed workflow show that SVR and GPR always provide the highest, or second-highest, predictive accuracy for the considered datasets. These results suggest that kernel-based models, and in particular distance-based similarity measures, are well suited to approximate the FEM response surfaces involved in this problem. This observation motivates the assessment of Radial Basis Function (RBF) interpolation⁴⁸

as a simpler and computationally efficient surrogate alternative. Although surrogate evaluation is negligible compared with FEM simulations, the training and hyperparameter optimization of complex supervised models can still become relevant when several outputs, geometry families and adaptive iterations are considered. RBF interpolation provides a non-parametric and fast-to-train alternative that can capture nonlinear response surfaces, making it attractive for low-dimensional and moderately sampled design spaces.

4.3 | RBF surrogate models

The RBF surrogate is implemented using the interpolator available in the *SciPy* library⁴⁹. For each output variable, the model is trained using the same input features employed by the supervised ML surrogates. In the current implementation, a multiquadric radial basis function is adopted with zero smoothing and automatic shape-parameter selection.

The prediction model can be expressed as

$$\hat{y}(\mathbf{x}) = \sum_{j=1}^N \lambda_j \phi(\|\mathbf{x} - \mathbf{x}_j\|), \quad (2)$$

where $\mathbf{x}_j$ denotes the FEM-sampled geometries, λ_j represents the interpolation weights and ϕ corresponds to the selected radial basis function.

For each output variable, a final RBF surrogate is trained using all available FEM samples in the current iteration and stored for subsequent use in the optimization process. Its predictive performance is assessed through Leave-One-Out validation: each FEM sample is iteratively removed from the dataset, a temporary RBF interpolant is trained with the remaining samples and the excluded sample is predicted. The resulting out-of-sample predictions are then used to compute RMSE, MAE and R^2 in order to assess the interpolation accuracy.

5 | Damage-aware surrogate-assisted optimization

The proposed methodology seeks to minimize local damage in both the dissipative windows and the surrounding frame, while promoting a balanced contribution of all windows to the energy dissipation process. The geometric optimization is carried out using DE⁵⁰, a population-based global optimizer that does not require gradient information and is therefore suitable for nonlinear and non-convex surrogate response surfaces. In the current implementation, DE is run with a maximum of 500 iterations, a population size factor of 25 and a convergence tolerance of 10^{-6} . Once an optimal candidate is obtained, an adaptive FEM validation loop is applied to verify the predicted geometry before acceptance.

For each candidate geometry $\mathbf{x}$, the trained surrogate models predict the window distortions $\hat{\varepsilon}_{xy,i}$, the window damage indicators $\hat{D}_i$, and the frame damage indicator $\hat{D}_f$. Damage is therefore controlled in all regions of the device, but with different mechanical relevance: frame damage is penalized more severely because it may compromise the structural integrity of the damper, whereas the window penalties are formulated to promote comparable damage levels among windows and avoid concentrating

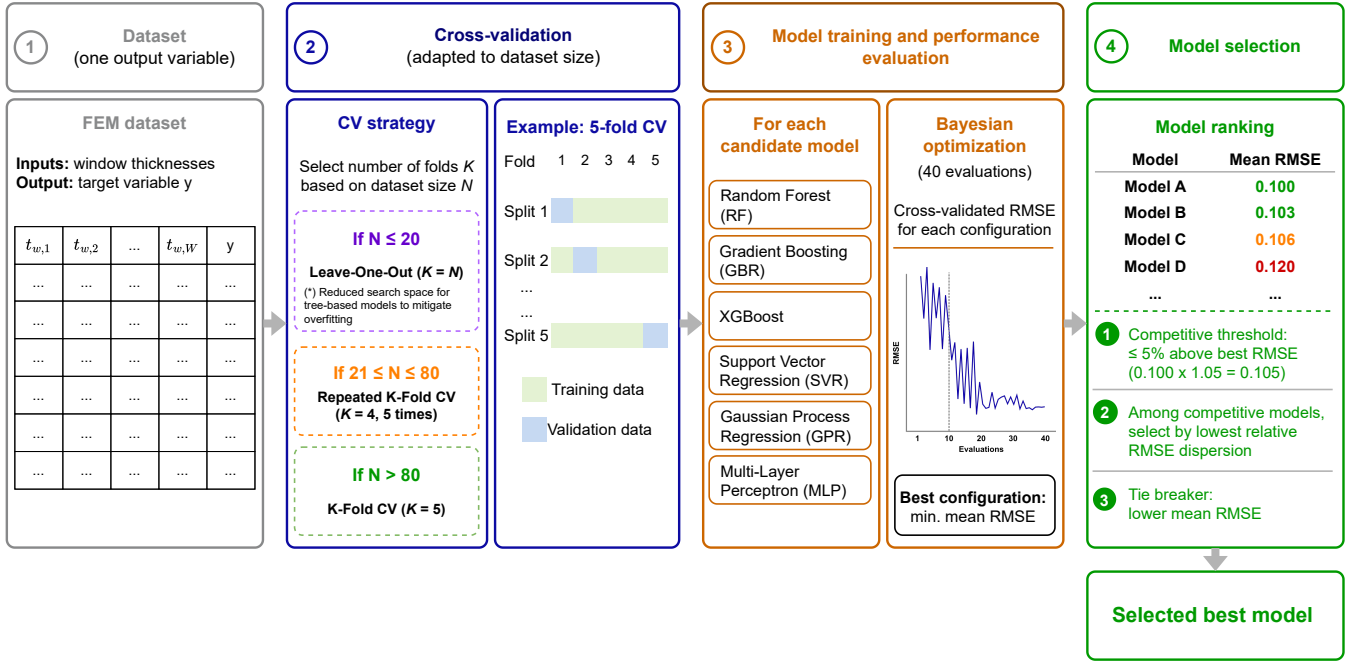


FIGURE 7 | Workflow of the supervised surrogate training and selection strategy. For each output variable, the cross-validation strategy is adapted to the dataset size, Bayesian optimization is used to tune each candidate model and the final surrogate is selected according to RMSE accuracy and fold-wise RMSE dispersion.

the dissipative demand in a single region. The dissipative contribution is estimated from $\hat{\varepsilon}_{xy,i}^2$, the window thickness and the corresponding area of the window, since the energy dissipated by each window depends not only on the distortion level but also on the amount of material involved. This term is several orders of magnitude smaller than the damage penalties and is intentionally left unscaled. As a result, damage control remains the dominant criterion, while the dissipative term acts as a tie-breaker among geometries with similar damage performance, favouring those with higher distortion and, consequently, greater energy dissipation capacity.

The implemented objective function to be minimized is

$$J(\mathbf{x}) = -\sum_{i=1}^W \hat{\varepsilon}_{xy,i}^2 t_{w,i} A_i + \sum_{i=1}^W P_w(\hat{\mathcal{D}}_i; \mathcal{D}_w^*) + P_f(\hat{\mathcal{D}}_f; \mathcal{D}_f^{\max}), \quad (3)$$

where W is the number of windows, $t_{w,i}$ is the thickness of window i , A_i is the corresponding area factor, $\mathcal{D}_w^*$ is the target damage level for the windows and $\mathcal{D}_f^{\max}$ is the maximum admissible frame damage threshold. The first term is negative because the optimizer minimizes J ; therefore, larger energy dissipation contributions reduce the objective value and are favoured once the damage criteria are satisfied.

The window penalty is defined as

$$P_w(\hat{\mathcal{D}}_i; \mathcal{D}_w^*) = \begin{cases} (\hat{\mathcal{D}}_i - \mathcal{D}_w^*)^2, & \hat{\mathcal{D}}_i > \mathcal{D}_w^*, \\ |\hat{\mathcal{D}}_i - \mathcal{D}_w^*|, & \hat{\mathcal{D}}_i \leq \mathcal{D}_w^*. \end{cases} \quad (4)$$

This formulation penalizes values above the target quadratically, while also discouraging excessively underused windows through a linear distance to the target. As a result, the optimizer tends to balance the damage levels among windows rather than forcing all windows to remain far below the admissible level.

The frame penalty is defined as

$$P_f(\hat{\mathcal{D}}_f; \mathcal{D}_f^{\max}) = \begin{cases} (\hat{\mathcal{D}}_f - \mathcal{D}_f^{\max})^3, & \hat{\mathcal{D}}_f > \mathcal{D}_f^{\max}, \\ 0, & \hat{\mathcal{D}}_f \leq \mathcal{D}_f^{\max}. \end{cases} \quad (5)$$

The cubic penalty reflects the greater severity of frame damage compared with localized damage in a window. Since a damage value $\mathcal{D} = 100$ represents complete failure, $\mathcal{D}_f^{\max} = 90$ is adopted in the current implementation. The target window threshold is also set to $\mathcal{D}_w^* = 90$.

The surrogate-optimized geometry is not accepted directly. Instead, once an optimal geometry is identified by the surrogate-assisted optimizer, it is re-evaluated with FEM to verify that the surrogate remains accurate in the region of the design space where the optimum lies. This validation step checks whether the surrogate has remained reliable in the region of the design space selected by the optimizer. The candidate geometry is accepted only if: (i) the prediction error of all damage and distortion variables remains below the prescribed tolerance, equal to 5%; (ii) the absolute error of the objective function remains within the admissible limit, set to 10; and (iii) the optimized window thicknesses remain stable between consecutive optimization iterations, with variations smaller than the 5% of the full design range. If all criteria are satisfied, the FEM-validated geometry is accepted as the optimized design. If at least one criterion is not satisfied, the new FEM result is added to the dataset, the surrogate models are retrained and the DE optimization is repeated. This loop, that reduces the risk of accepting a geometry that is optimal only because of surrogate extrapolation error, is summarized in Figure 8.

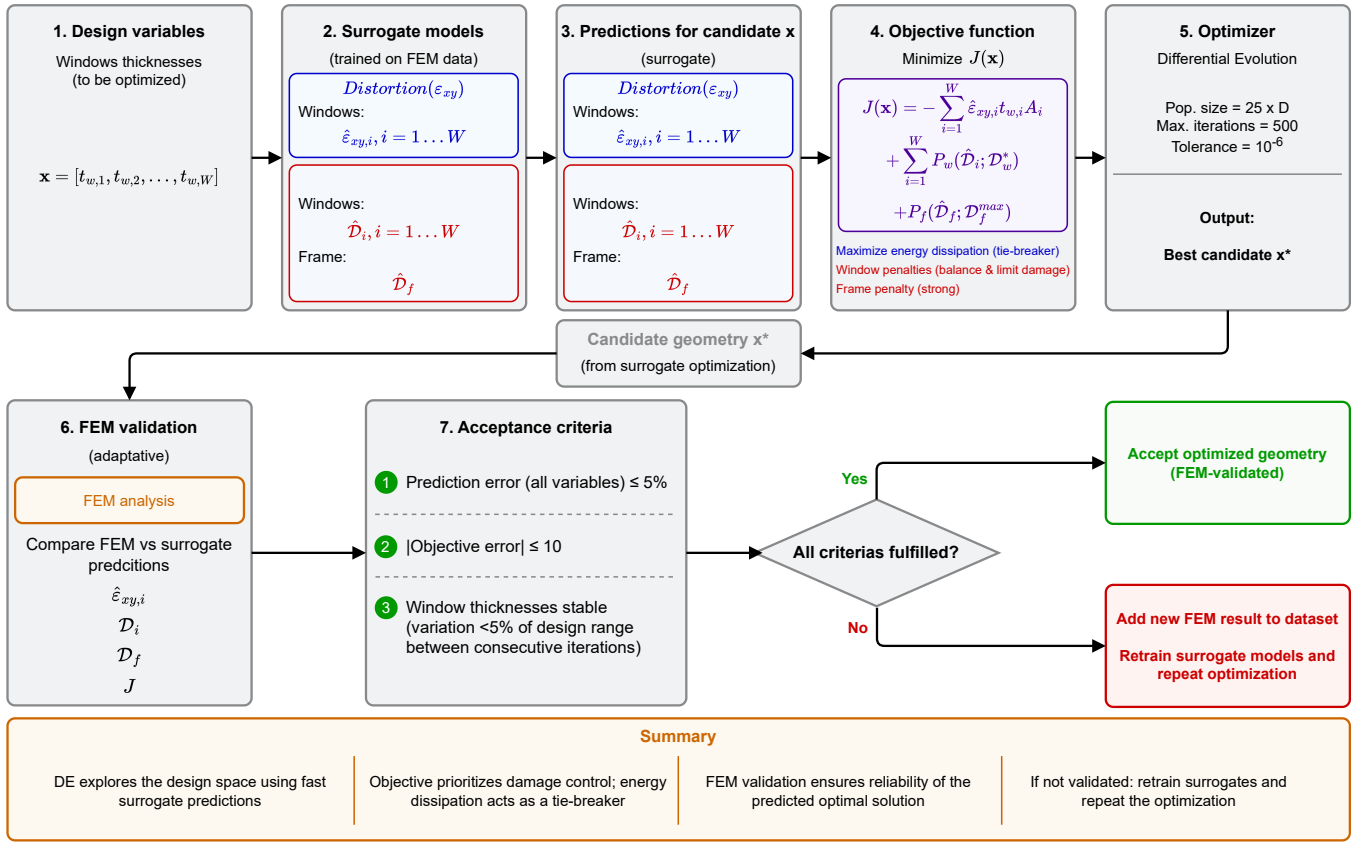


FIGURE 8 | Surrogate-assisted optimization and FEM validation retraining loop.

6 | Numerical results and discussion

The supervised-learning comparison shows a clear hierarchy among the candidate surrogate models. A total of 100 output-specific training problems were considered, corresponding to the $2W + 1$ target variables required for each geometry family and adaptive iteration. The F_1 family required two optimization iterations, whereas the remaining families required three. Across all geometry families, iterations and output variables, SVR was the most frequently selected model and also the model that most often achieved the lowest cross-validated RMSE. As summarized in Figure 9, SVR was selected in 71 cases, followed by GPR, GBR, XGBoost and MLP, while Random Forest was not selected in any case. This dominance was particularly clear for the two-window families, where only one output was assigned to a GPR model and for the damage-related outputs, for which SVR was selected in 47 out of 57 cases. For more complex devices, those with three and five windows, and for distortion-related outputs, the model selection became more heterogeneous, although SVR still provided the best overall performance. From the computational point of view, SVR also offered the lowest median training times among the supervised models, whereas MLP required substantially longer training times without providing gain in accuracy.

These results indicate that kernel-based models are particularly well suited to the present surrogate task. SVR provides the best compromise between accuracy and computational cost, while GPR is the second most competitive supervised strategy, especially in some higher-dimensional cases. Tree-based models, although robust, are less frequently selected and MLP models

are not competitive in terms of computational efficiency for the dataset sizes considered here. As mentioned in previous sections, this behaviour motivated the additional evaluation of RBF interpolation as a simpler surrogate alternative. In contrast to the supervised models, RBF models were trained in less than one second per output, making them especially attractive for repeated surrogate updates within the adaptive optimization loop.

The FEM validation of the optimized geometries is summarized in Table 3, while the complete optimization results for all adaptive iterations are provided in Appendix B. For each geometry family and surrogate strategy, Table 3 reports the final accepted adaptive iteration, the optimized window thicknesses t_w^* , the surrogate-predicted objective value J_{sur} , the corresponding FEM-recomputed objective value J_{FEM} , and the associated validation errors, $|e_J|$ and e_{max} . The maximum variable error, e_{max} , is defined as the largest relative error among all quantities entering the objective function, namely $\varepsilon_{xy,i}$, $\mathcal{D}_i$ and $\mathcal{D}_f$, whereas $|e_J|$ denotes the absolute objective-function error. The optimization required between two and three adaptive iterations depending on the geometry family and surrogate type, with most cases converging after three iterations. No systematic difference in the number of iterations was observed between RBF and supervised ML surrogates.

The final validation results show that both surrogate strategies provide FEM-consistent optimized geometries after only a few adaptive iterations. RBF surrogates generally lead to lower objective-function discrepancies, with an average $|e_J|$ of 1.36, compared with 5.04 for the supervised ML surrogates. The average maximum variable error is also lower for RBF, with 1.79%




FIGURE 9 | Summary of the supervised surrogate selection over all geometry families, adaptive iterations and target outputs. The training time corresponds to the median time required for one Bayesian hyperparameter search for a single output variable and, for visualization purposes, is in logarithmic scale.

TABLE 3 | Final FEM validation of the optimized geometries obtained with supervised ML and RBF surrogates.

Family	Surrogate	Iterations	t_w^* [mm]	J_{surr}	J_{FEM}	$ e_J $	e_{max} [%]
F_1	RBF	3	[12.56, 14.79]	0.00	0.41	0.41	0.29
F_1	Supervised ML	2	[12.53, 14.75]	0.00	0.08	0.08	3.14
F_2	RBF	3	[14.77, 18.95]	688.53	691.06	2.53	0.10
F_2	Supervised ML	3	[15.20, 20.00]	796.29	802.64	6.35	0.64
F_3	RBF	3	[5.81, 7.88, 8.98]	0.00	0.93	0.93	1.70
F_3	Supervised ML	3	[5.72, 7.87, 8.96]	0.00	7.47	7.47	2.85
F_4	RBF	2	[6.81, 9.02, 9.65]	0.00	1.39	1.39	4.21
F_4	Supervised ML	3	[6.84, 9.05, 9.73]	0.00	1.35	1.35	2.91
F_5	RBF	3	[5.77, 7.38, 8.37, 6.18, 5.00]	17.74	19.26	1.51	2.66
F_5	Supervised ML	3	[5.74, 7.46, 8.50, 6.37, 5.00]	16.90	26.87	9.97	3.81

compared with 2.67% for supervised ML. The RBF surrogate provides particularly accurate predictions for the two-window devices, with e_{max} below 0.3%, while remaining below 4.3% for the three- and five-window families. The supervised ML surrogates also satisfy all validation criteria, although larger discrepancies are observed in some cases, especially for the F_5 family, where the objective-function error reaches 9.97.

The adaptive validation loop is essential to reach these levels of agreement. In several cases, the first surrogate-optimized candidate did not satisfy the prescribed error tolerances, particularly for the three- and five-window devices. After incorporating the additional FEM results and retraining the surrogates, the prediction errors decreased significantly. For instance, the maximum variable error of the RBF surrogate decreased from 13.70% to 1.70% in the F_3 family and from 11.44% to 2.66% in the F_5 family. A similar behaviour was observed for the supervised ML surrogates, whose final candidates also satisfied all acceptance criteria. This confirms that the adaptive loop reduces the risk of accepting geometries that appear optimal only because of surrogate prediction errors in sparsely sampled regions of the design space.

The optimized geometries obtained with RBF and supervised ML surrogates are very similar in most geometry families. In general,

the differences in the optimized window thicknesses are below 0.1 mm and always below 0.2 mm, except for the F_2 family. In this case, the differences reach 0.43 mm and 1.05 mm for the first and second windows, respectively. This family also presents the largest objective-function values, indicating that the prescribed damage targets, $\mathcal{D}_i$ and $\mathcal{D}_f$, cannot be fully maintained below the desired threshold within the admissible thickness range. Consequently, the optimizer is forced to find the best compromise between damage control and dissipative activation, which may amplify the differences between surrogate predictions.

Figure 10 shows the evolution of the window thicknesses during the adaptive optimization process. In most cases, the largest adjustment occurs between the first and second iterations, while the changes between the second and third iterations are considerably smaller, indicating convergence of the proposed geometry. According to the acceptance criteria, the maximum variation in any window thickness between the last two iterations must remain below 5% of the corresponding design range. This condition is satisfied in all validated cases, with a maximum observed variation of 4.44%. Therefore, the adaptive process converges not only in terms of surrogate prediction accuracy, but also in terms of the optimized geometry proposed by the optimizer.




FIGURE 10 | Evolution of the optimized window thicknesses during the adaptive optimization process.

From a methodological point of view, these results highlight the trade-off between surrogate complexity, accuracy and computational efficiency. Supervised models, particularly SVR and GPR, provide high predictive accuracy and robustness, but require hyperparameter optimization and cross-validation for every output variable and adaptive iteration. RBF interpolation, in contrast, has a much lower training cost and provides very competitive final predictions once the relevant regions of the design domain have been adaptively sampled. Therefore, the comparison does not identify a universally superior surrogate strategy. Instead, it suggests that RBF interpolation is especially suitable for low- to moderate-dimensional design spaces with well-distributed FEM samples and relatively smooth input-output relationships, as occurs for the response variables analysed in this work. Supervised ML surrogates remain valuable when greater robustness is required, or when the response surface is expected to involve stronger nonlinear interactions, local irregularities or higher-dimensional dependencies.

The objective-function values should also be interpreted carefully. The objective-function error measures the consistency between surrogate predictions and FEM validation, not whether the final objective value is necessarily close to zero. Some geometry families, such as F_2 , retain non-negligible penalty contributions because the prescribed damage targets cannot be fully achieved within the admissible design bounds. Nevertheless, the close agreement between surrogate and FEM objective values indicates that the accepted designs are not artifacts of surrogate extrapolation, but FEM-consistent optimized candidates within the explored design space.

This behaviour is illustrated in Figure 11, which shows the evolution of the RBF objective surface during the adaptive optimization process for the two-window families. The response surfaces remain relatively smooth, supporting the suitability of RBF interpolation for this problem. The evolution between adaptive iterations also shows how the surrogate surface is progressively

corrected as new FEM information is incorporated near the optimized region.

7 | Conclusions and future work

This work presented an adaptive surrogate-assisted optimization framework for buckling-delayed shear-link dampers subjected to cyclic seismic loading. The proposed methodology addresses the main limitation of direct FEM-based optimization, the high computational cost associated with repeatedly evaluating nonlinear cyclic simulations. By training surrogate models on FEM-generated datasets and validating the optimized candidates through additional FEM analyses, the optimizer can efficiently explore the design domain while remaining consistent with the mechanical response captured by the calibrated numerical model.

One of the key features of the proposed optimization framework, compared to approaches focused primarily on energy maximization, is the formulation of an objective function that takes damage into account. The aim is to prioritise local damage control in both the dissipative windows and the surrounding frame, whilst promoting a balanced contribution from all windows to the energy dissipation process. Thus, damage to the windows is permitted and expected, as they are intended to act as dissipative regions, provided that it remains controlled and reasonably distributed; conversely, damage to the frame is penalised more severely because it can compromise the structural integrity of the damper. Therefore, rather than merely maximizing the energy dissipated, the proposed formulation favours geometries that concentrate dissipative activation in the windows, prevent excessive damage localization in a single region and protect the frame from critical damage.

A comprehensive comparison of surrogate strategies was performed in terms of predictive accuracy and computational efficiency. The supervised ML results showed that SVR and GPR




FIGURE 11 | Evolution of the RBF objective surface during the adaptive optimization process for the two-window families. Left: F_1 . Right: F_2 .

were the most competitive models, with SVR being the most frequently selected across the analysed outputs. RBF interpolation proved to be even a better alternative, due to its high efficiency: in the final adaptive iterations, it achieved validation errors comparable to, and in most cases lower than, those of the supervised ML surrogates, while requiring substantially lower training effort and no hyperparameter search. This performance is attributed to the characteristics of the present problem: low- to moderate-dimensional design spaces, well-distributed FEM samples and nonlinear but relatively smooth relationships between window thicknesses and response indicators.

The proposed adaptive validation loop proved to be necessary and effective. Several initially optimized candidates did not satisfy the prescribed error tolerances. After incorporating the new FEM results into the training dataset and retraining the surrogates, the prediction errors decreased and all final optimized geometries satisfied the acceptance criteria after only two or three iterations. Therefore, the final designs are not accepted solely on the basis of surrogate predictions, but are explicitly verified through FEM in the region of the design space where the optimum is located.

The proposed methodology also has some limitations that should be acknowledged. First, its reliability depends on the quality of the calibrated FEM model used to generate the training data and validate the optimized designs. Second, the TFDMap is used here as a post-processing damage indicator rather than as a constitutive fracture model; therefore, the optimized configurations should be interpreted in terms of relative damage control and proximity to critical states, not as direct predictions of crack initiation. Third, only the window thicknesses are considered as design variables. Although this leads to a controlled and interpretable optimization problem, it does not exploit the full geometric flexibility of BDSL dampers. Finally, the optimized geometries should ultimately be validated experimentally before being used to establish general design recommendations.

Future work should extend the design space by including additional geometric and mechanical variables, such as window height, window spacing, frame thickness or global device proportions. This extension would increase the dimensionality and

complexity of the surrogate task. In those cases, the performance of RBF interpolation should therefore be reassessed. While RBF models performed very well in the present study, their efficiency and accuracy may decrease as the input space becomes larger or the response surfaces develop stronger local nonlinearities. In such cases, supervised ML models or hybrid surrogate strategies may become more advantageous.

Another relevant future direction is the incorporation of interpretability analyses, such as SHapley Additive exPlanations, to quantify the influence of each geometric variable on window damage, frame damage and dissipative activation. Although such analysis lies outside the main scope of the present study, it could provide valuable insight into the design drivers governing the behaviour of BDSL dampers and support more transparent engineering decision-making. Overall, the proposed methodology establishes a scalable basis for FEM-consistent, damage-aware optimization of seismic energy dissipation devices, while leaving room for broader design variables, richer surrogate strategies and experimental validation of the optimized configurations.

AUTHOR CONTRIBUTIONS

Conceptualization: J. Irazábal, J. Ramírez, J. M. González, G. Bozzo, L. Bozzo. Methodology: J. Irazábal, J. Ramírez, J. M. González. Software and surrogate optimization: J. Irazábal. FEM modelling and validation: J. Ramírez, J. M. González, L. Lázaro, F. Rastellini. Writing—original draft: J. Irazábal. Writing—review and editing: all authors.

ACKNOWLEDGMENTS

The authors acknowledge the financial support of Project ACE100/23/000022, “Edificacions resilientes equipades amb dissipadores Shear Link”, funded by the Government of Catalonia through ACCIO and with the support of the Catalan Office for Climate Change, with the participation of Luis Bozzo Estructuras y Proyectos S.L. and the Centre Internacional de Metodologies Numerics en Enginyeria (CIMNE).

FINANCIAL DISCLOSURE

None reported.

CONFLICT OF INTEREST

The authors declare no potential conflict of interests.

References

1. Malley JO, Popov EP. Shear Links in Eccentrically Braced Frames. *Journal of Structural Engineering*. 1984;110(9):2275–2295. doi: 10.1061/(ASCE)0733-9445(1984)110:9(2275)
2. Okazaki T, Engelhardt MD. Cyclic loading behavior of EBF links constructed of ASTM A992 steel. *Journal of Constructional Steel Research*. 2007;63(6):751–765. doi: 10.1016/j.jcsr.2006.08.004
3. Deng K, Pan P, Su Y, Ran T, Xue Y. Development of an energy dissipation restrainer for bridges using a steel shear panel. *Journal of Constructional Steel Research*. 2014;101:83–95. doi: 10.1016/j.jcsr.2014.03.009
4. Deng K, Pan P, Li W, Xue Y. Development of a buckling restrained shear panel damper. *Journal of Constructional Steel Research*. 2015;106:311–321. doi: 10.1016/j.jcsr.2015.01.004
5. Deng K, Pan P, Sun J, Liu J, Xue Y. Shape optimization design of steel shear panel dampers. *Journal of Constructional Steel Research*. 2014;99:187–193. doi: 10.1016/j.jcsr.2014.03.001
6. Deng K, Pan P, Su Y, Xue Y. Shape optimization of U-shaped damper for improving its bi-directional performance under cyclic loading. *Engineering Structures*. 2015;93:27–35. doi: 10.1016/j.engstruct.2015.03.006
7. Motamedi M, Nateghi-A. F. Study on mechanical characteristics of accordion metallic damper. *Journal of Constructional Steel Research*. 2018;142:68–77. doi: 10.1016/j.jcsr.2017.12.010
8. Ghamari A, Kim YJ, Bae J. Utilizing an I-shaped shear link as a damper to improve the behaviour of a concentrically braced frame. *Journal of Constructional Steel Research*. 2021;186:106915. doi: 10.1016/j.jcsr.2021.106915
9. Xiong L, Guo Z, Cai J, Jiang K, Li L. Seismic performance of the replaceable steel links with different short length ratios. *Sci Rep*. 2024;14(1):29976. doi: 10.1038/s41598-024-81336-8
10. Deng K, Pan P, Lam A, Xue Y. A simplified model for analysis of high-rise buildings equipped with hysteresis damped outriggers. *The Structural Design of Tall and Special Buildings*. 2014;23(15):1158–1170. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/tal.1113>doi: 10.1002/tal.1113
11. Kim YC, Mortazavi SJ, Farzampour A, Hu JW, Mansouri I, Awoyera PO. Optimization of the Curved Metal Damper to Improve Structural Energy Dissipation Capacity. *Buildings*. 2022;12(1):67. doi: 10.3390/buildings12010067
12. Zhang Z, Ou J, Li D, Zhang S. Optimization Design of Coupling Beam Metal Damper in Shear Wall Structures. *Applied Sciences*. 2017;7(2):137. doi: 10.3390/app7020137
13. Farzampour A, Khatibinia M, Mansouri I. Shape optimization of butterfly-shaped shear links using Grey Wolf algorithm. *Ingegneria Sismica*. 2019;36:27–41.
14. Khatibinia M, Jalaipour M, Gharehbaghi S. Shape optimization of U-shaped steel dampers subjected to cyclic loading using an efficient hybrid approach. *Engineering Structures*. 2019;197:108874. doi: 10.1016/j.engstruct.2019.02.005
15. Khatibinia M, Ahrari A, Gharehbaghi S, et al. An efficient approach for optimum shape design of steel shear panel dampers under cyclic loading. *Smart Structures and Systems*. 2021;27(3):547–557. doi: 10.12989/sss.2021.27.3.547
16. Shi JX, Kozono S, Shimoda M, Takino M, Wada D, Liu Y. Non-parametric shape design optimization of elastic-plastic shear panel dampers under cyclic loading. *Engineering Structures*. 2019;189:48–61. doi: 10.1016/j.engstruct.2019.03.049
17. Saleh YN, Mourad SA, Ibrahim AM. Topology optimization of vertical shear links in eccentrically braced frames. *Structures*. 2024;66:106821. doi: 10.1016/j.istruc.2024.106821
18. Saleh YN, Mourad SA, Salem HG, Ibrahim AM. Computational study on stiffened topology-optimized shear links for eccentrically braced frames. *Bull Earthquake Eng*. 2026. doi: 10.1007/s10518-026-02417-9
19. Mendoza-Cuy A, Begambre-Carrillo O, Villalba-Morales JD. Topology optimization of steel slotted dampers with the hybrid cellular automata technique. *Advances in Engineering Software*. 2025;206:103921. doi: 10.1016/j.advensoft.2025.103921
20. Ríos I, Gómez Á, Romero F, et al. Computational Design of High-Performance U-Shaped Seismic Dampers Using Statistical Optimization. *Materials*. 2025;18(23):5403. doi: 10.3390/ma18235403
21. Chan RWK, Yuen JKK, Lee EWM, Arashpour M. Application of Nonlinear-Autoregressive-Exogenous model to predict the hysteretic behaviour of passive control systems. *Engineering Structures*. 2015;85:1–10. doi: 10.1016/j.engstruct.2014.12.007
22. Bae J, Lee CH, Park M, Alemayehu RW, Ryu J, Ju YK. Modified Low-Cycle Fatigue Estimation Using Machine Learning for Radius-Cut Coke-Shaped Metallic Damper Subjected to Cyclic Loading. *Int J Steel Struct*. 2020;20(6):1849–1858. doi: 10.1007/s13296-020-00377-7
23. Almasabha G, Alshboul O, Shehadeh A, Almuflhi AS. Machine Learning Algorithm for Shear Strength Prediction of Short Links for Steel Buildings. *Buildings*. 2022;12(6):775. doi: 10.3390/buildings12060775
24. Elgammal A, Ali Y. A novel hysteretic restoring force model for shear link dampers: A machine learning approach. *Structures*. 2024;70:107848. doi: 10.1016/j.istruc.2024.107848
25. Hu S, Wang W, Lu Y. Explainable machine learning models for probabilistic buckling stress prediction of steel shear panel dampers. *Engineering Structures*. 2023;288:116235. doi: 10.1016/j.engstruct.2023.116235
26. Hu Y, Guo W, Long Y, Li S, Xu Z. Physics-informed deep neural networks for simulating S-shaped steel dampers. *Computers & Structures*. 2022;267:106798. doi: 10.1016/j.compstruc.2022.106798
27. Ramírez Machado BJ, Bozzo Fernández G, González Lopez JM, Rastellini FG, Irazábal González J, Bozzo Rotondo LM. Advanced modeling of buckling delayed shear links under strong cyclic loads. In: National Technical University of Athens. 2025.
28. Martínez X, Rastellini F, Oller S, Flores F, Oñate E. Computationally optimized formulation for the simulation of composite materials and delamination failures. *Composites Part B: Engineering*. 2011;42(2):134–144. doi: 10.1016/j.compositesb.2010.09.013
29. Yoshida F, Uemori T. A model of large-strain cyclic plasticity describing the Bauschinger effect and workhardening stagnation. *International Journal of Plasticity*. 2002;18(5–6):661–686. doi: 10.1016/S0749-6419(01)00050-X
30. Jia LJ. Integration algorithm for a modified Yoshida–Uemori model to simulate cyclic plasticity in extremely large plastic strain ranges up to fracture. *Computers & Structures*. 2014;145:36–46. doi: 10.1016/j.compstruc.2014.08.010
31. Seismic provisions for structural steel buildings. Standard, American Institute of Steel Construction; 2002.
32. Minimum design loads and associated criteria for buildings and other structures. Standard, American Society of Civil Engineers; 2017.
33. Rastellini F, Socorro G, Forgas A, Oñate E. A Triaxial Failure Diagram to predict the forming limit of 3D sheet metal parts subjected to multiaxial stresses. *J. Phys.: Conf. Ser.*. 2016;734(3):032020. doi: 10.1088/1742-6596/734/3/032020
34. Rice JR, Tracey DM. On the ductile enlargement of voids in triaxial stress fields. *Journal of the Mechanics and Physics of Solids*. 1969;17(3):201–217. doi: 10.1016/0022-5096(69)90033-7
35. Bao Y, Wierzbicki T. On fracture locus in the equivalent strain and stress triaxiality space. *International Journal of Mechanical Sciences*. 2004;46(1):81–98. doi: 10.1016/j.ijmecsci.2004.02.006
36. Wierzbicki T, Bao Y, Lee YW, Bai Y. Calibration and evaluation of seven fracture models. *International Journal of Mechanical*

- Sciences*. 2005;47(4):719–743. doi: [10.1016/j.ijmecsci.2005.03.003](https://doi.org/10.1016/j.ijmecsci.2005.03.003)
37. Bai Y, Wierzbicki T. A new model of metal plasticity and fracture with pressure and Lode dependence. *International Journal of Plasticity*. 2008;24(6):1071–1096. doi: [10.1016/j.ijplas.2007.09.004](https://doi.org/10.1016/j.ijplas.2007.09.004)
 38. Joseph VR, Hung Y. Orthogonal-Maximin Latin Hypercube Designs. *Statistica Sinica*. 2008;18(1):171–186.
 39. Sheikholeslami R, Razavi S. Progressive Latin Hypercube Sampling: An efficient approach for robust sampling-based analysis of environmental models. *Environmental Modelling & Software*. 2017;93:109–126. doi: [10.1016/j.envsoft.2017.03.010](https://doi.org/10.1016/j.envsoft.2017.03.010)
 40. Breiman L. Random Forests. *Machine Learning*. 2001;45(1):5–32. doi: [10.1023/A:1010933404324](https://doi.org/10.1023/A:1010933404324)
 41. Friedman J. Greedy function approximation: a gradient boosting machine. *Annals of Statistics*. 2001:1189 – 1232.
 42. Chen T, Guestrin C. XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 2016:785–794. Conference Name: KDD '16: The 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining ISBN: 9781450342322doi: [10.1145/2939672.2939785](https://doi.org/10.1145/2939672.2939785)
 43. Drucker H, Burges CJC, Kaufman L, Smola A, Vapnik V. Support vector regression machines. In: *NIPS'96*. MIT Press. 1996; Cambridge, MA, USA:155–161.
 44. Williams CKI, Rasmussen CE. Gaussian processes for regression. In: *NIPS'95*. MIT Press. 1995; Cambridge, MA, USA:514–520.
 45. Rosenblatt F. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review*. 1958;65(6):386–408. doi: [10.1037/h0042519](https://doi.org/10.1037/h0042519)
 46. Rumelhart DE, Hinton GE, Williams RJ. Learning representations by back-propagating errors. *Nature*. 1986;323(6088):533–536. doi: [10.1038/323533a0](https://doi.org/10.1038/323533a0)
 47. Snoek J, Larochelle H, Adams RP. Practical Bayesian optimization of machine learning algorithms. In: . 2 of *NIPS'12*. Curran Associates Inc. 2012; Red Hook, NY, USA:2951–2959
 48. Gutmann HM. A Radial Basis Function Method for Global Optimization. *Journal of Global Optimization*. 2001;19(3):201–227. doi: [10.1023/A:1011255519438](https://doi.org/10.1023/A:1011255519438)
 49. Virtanen P, Gommers R, Oliphant TE, et al. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*. 2020;17:261–272. doi: [10.1038/s41592-019-0686-2](https://doi.org/10.1038/s41592-019-0686-2)
 50. Storn R, Price K. Differential Evolution – A Simple and Efficient Heuristic for global Optimization over Continuous Spaces. *Journal of Global Optimization*. 1997;11(4):341–359. doi: [10.1023/A:1008202821328](https://doi.org/10.1023/A:1008202821328)

SUPPORTING INFORMATION

Additional information related to this study is available from the corresponding author upon reasonable request.

APPENDIX

A | Hyperparameter search spaces for supervised ML surrogate models

Table A1 summarizes the hyperparameter search spaces used for the Bayesian optimization of the supervised surrogate models. The same ranges are used for kernel-based and neural-network models across all dataset sizes, whereas the search spaces of tree-based models are automatically restricted for small datasets to reduce overfitting. The cross-validation strategy is also adapted to the dataset size: Leave-One-Out validation is used for $N \leq 20$, repeated four-fold cross-validation with five repetitions for $21 \leq N \leq 80$, and shuffled five-fold cross-validation for larger datasets.

B | Summary of optimization results

This appendix summarizes the surrogate-predicted optimized configurations, the corresponding FEM validation values and the associated errors for all adaptive iterations. The source “Surr.” denotes the values predicted by the surrogate model during the optimization, whereas “FEM” denotes the values recomputed with the high-fidelity numerical model. Error rows report the relative error for the response variables and the absolute error for the objective function.

TABLE A1 | Hyperparameter search spaces used for Bayesian optimization of the supervised surrogate models.

Model	Preprocessing / kernel	Hyperparameter	Search space
RF	–	$n_{\text{estimators}}$	[200, 2000]
		max depth	[5, 40]
		min samples split	[2, 10]
		min samples leaf	[1, 6]
		max features	[0.5, 1.0]
GBR	–	$n_{\text{estimators}}$	[200, 3000]
		learning rate	$[10^{-3}, 10^{-1}]$
		max depth	[2, 6]
		subsample	[0.6, 1.0]
		max features	[0.5, 1.0]
XGBoost	–	$n_{\text{estimators}}$	[200, 2000]
		max depth	[3, 8]
		learning rate	$[10^{-3}, 0.3]$
		subsample	[0.6, 1.0]
		colsample by tree	[0.6, 1.0]
		min child weight	[1, 10]
SVR	StandardScaler + RBF kernel	C	$[10^0, 10^4]$
		γ	$[10^{-4}, 10^{-1}]$
		ϵ	$[10^{-4}, 10^{-1}]$
GPR	StandardScaler	kernel	{RBF, Matern _{3/2} , Matern _{5/2} , RQ}
		amplitude	$[10^{-2}, 10^2]$
		length scale	$[10^{-2}, 10^2]$
		α_{RQ}	$[10^{-2}, 10^2]$
		noise	$[10^{-12}, 10^{-4}]$
		restarts	[1, 10]
MLP	StandardScaler	hidden layers	[1, 5]
		neurons per layer	[64, 512]
		activation	{ReLU, tanh}
		α	$[10^{-5}, 10^{-2}]$
		initial learning rate	$[10^{-4}, 10^{-2}]$

TABLE B2 | Optimization results for the F_1 family using Supervised ML surrogates.

Iter.	Source	$\mathbf{t}_w$ [mm]	ϵ_{xy}	$\mathcal{D}_w$	$\mathcal{D}_f$	J
1	Surr.	[12.60, 14.64]	[0.05048, 0.06183]	[90.00, 90.00]	89.65	0.00
	FEM	[12.60, 14.64]	[0.04501, 0.0616]	[88.25, 91.19]	87.17	3.16
	Error	–	[12.15%, 0.37%]	[1.99%, 1.30%]	2.84%	3.16
2	Surr.	[12.53, 14.75]	[0.04602, 0.06061]	[90.00, 90.00]	89.32	0.00
	FEM	[12.53, 14.75]	[0.04555, 0.06053]	[89.95, 90.18]	86.61	0.07947
	Error	–	[1.03%, 0.12%]	[0.05%, 0.20%]	3.14%	0.0795

TABLE B3 | Optimization results for the F_2 family using Supervised ML surrogates.

Iter.	Source	$\mathbf{t}_w$ [mm]	ϵ_{xy}	$\mathcal{D}_w$	$\mathcal{D}_f$	J
1	Surr.	[15.69, 20.00]	[0.05065, 0.04732]	[105.29, 72.51]	96.62	541.95
	FEM	[15.69, 20.00]	[0.05127, 0.04603]	[103.51, 73.24]	99.21	980.46
	Error	–	[1.21%, 2.80%]	[1.72%, 0.99%]	2.61%	438.51
2	Surr.	[15.14, 20.00]	[0.05444, 0.04442]	[110.55, 71.35]	97.30	830.19
	FEM	[15.14, 20.00]	[0.05413, 0.04418]	[110.40, 71.84]	97.12	795.11
	Error	–	[0.58%, 0.53%]	[0.14%, 0.68%]	0.19%	35.08
3	Surr.	[15.20, 20.00]	[0.05398, 0.0441]	[109.73, 71.61]	97.30	796.29
	FEM	[15.20, 20.00]	[0.05379, 0.04439]	[109.65, 72.03]	97.36	802.64
	Error	–	[0.36%, 0.64%]	[0.07%, 0.59%]	0.06%	6.35

TABLE B4 | Optimization results for the F_3 family using Supervised ML surrogates.

Iter.	Source	$\mathbf{t}_w$ [mm]	ε_{xy}	$\mathcal{D}_w$	$\mathcal{D}_f$	J
1	Surr.	[5.96, 8.23, 9.47]	[0.05164, 0.04547, 0.04887]	[90.00, 90.00, 90.00]	71.00	0.00
	FEM	[5.96, 8.23, 9.47]	[0.04637, 0.04648, 0.05136]	[89.20, 85.22, 83.96]	68.08	11.62
	Error	–	[11.37%, 2.18%, 4.85%]	[0.90%, 5.61%, 7.19%]	4.28%	11.62
2	Surr.	[5.70, 7.90, 9.02]	[0.04801, 0.0478, 0.0493]	[90.00, 90.00, 90.00]	71.22	0.00
	FEM	[5.70, 7.90, 9.02]	[0.04906, 0.04815, 0.05333]	[93.70, 89.65, 88.89]	68.18	15.15
	Error	–	[2.14%, 0.73%, 7.55%]	[3.95%, 0.39%, 1.25%]	4.46%	15.15
3	Surr.	[5.72, 7.87, 8.96]	[0.04862, 0.04822, 0.05305]	[90.00, 90.00, 90.00]	68.58	0.00
	FEM	[5.72, 7.87, 8.96]	[0.04846, 0.04798, 0.05389]	[92.64, 89.69, 89.81]	68.52	7.47
	Error	–	[0.32%, 0.50%, 1.57%]	[2.85%, 0.34%, 0.21%]	0.09%	7.47

TABLE B5 | Optimization results for the F_4 family using Supervised ML surrogates.

Iter.	Source	$\mathbf{t}_w$ [mm]	ε_{xy}	$\mathcal{D}_w$	$\mathcal{D}_f$	J
1	Surr.	[7.34, 9.28, 10.13]	[0.03217, 0.03708, 0.05097]	[90.00, 90.00, 90.00]	71.57	0.00
	FEM	[7.34, 9.28, 10.13]	[0.03834, 0.04702, 0.05298]	[80.76, 86.78, 87.08]	69.66	15.38
	Error	–	[16.10%, 21.13%, 3.79%]	[11.44%, 3.71%, 3.35%]	2.75%	15.38
2	Surr.	[6.87, 9.08, 9.86]	[0.03612, 0.05013, 0.05362]	[90.00, 90.00, 90.00]	68.42	0.00
	FEM	[6.87, 9.08, 9.86]	[0.04277, 0.04322, 0.05375]	[90.40, 89.46, 88.14]	68.49	2.56
	Error	–	[15.56%, 15.99%, 0.24%]	[0.44%, 0.60%, 2.11%]	0.10%	2.56
3	Surr.	[6.84, 9.05, 9.73]	[0.04169, 0.04354, 0.05319]	[90.00, 90.00, 90.00]	68.81	0.00
	FEM	[6.84, 9.05, 9.73]	[0.04271, 0.04299, 0.05478]	[90.36, 89.19, 89.59]	68.94	1.35
	Error	–	[2.40%, 1.27%, 2.91%]	[0.40%, 0.90%, 0.46%]	0.19%	1.35

TABLE B6 | Optimization results for the F_5 family using Supervised ML surrogates.

Iter.	Source	$\mathbf{t}_w$ [mm]	ε_{xy}	$\mathcal{D}_w$	$\mathcal{D}_f$	J
1	Surr.	[5.97, 7.38, 8.56, 6.70, 5.00]	[0.05963, 0.05258, 0.06459, 0.06971, 0.05512]	[90.00, 90.00, 90.00, 90.00, 79.42]	74.65	10.58
	FEM	[5.97, 7.38, 8.56, 6.70, 5.00]	[0.05635, 0.05188, 0.05966, 0.07124, 0.05824]	[84.95, 90.65, 87.91, 81.06, 74.18]	79.15	32.33
	Error	–	[5.83%, 1.36%, 8.27%, 2.14%, 5.35%]	[5.94%, 0.72%, 2.38%, 11.03%, 7.07%]	5.69%	21.74
2	Surr.	[5.76, 7.37, 8.46, 6.67, 5.00]	[0.0615, 0.05313, 0.06014, 0.07067, 0.05734]	[90.00, 90.00, 90.00, 90.00, 73.04]	74.85	16.96
	FEM	[5.76, 7.37, 8.46, 6.67, 5.00]	[0.06398, 0.05679, 0.06492, 0.07844, 0.06236]	[95.56, 95.92, 92.86, 85.17, 71.37]	75.19	97.55
	Error	–	[3.88%, 6.45%, 7.36%, 9.90%, 8.05%]	[5.81%, 6.17%, 3.08%, 5.67%, 2.34%]	0.45%	80.59
3	Surr.	[5.74, 7.46, 8.50, 6.37, 5.00]	[0.06117, 0.05325, 0.0609, 0.07548, 0.05962]	[90.00, 90.00, 90.00, 90.00, 73.10]	75.62	16.90
	FEM	[5.74, 7.46, 8.50, 6.37, 5.00]	[0.06082, 0.05335, 0.05924, 0.07453, 0.05842]	[91.93, 89.75, 87.52, 86.70, 72.88]	78.19	26.87
	Error	–	[0.59%, 0.18%, 2.80%, 1.28%, 2.06%]	[2.10%, 0.28%, 2.83%, 3.81%, 0.30%]	3.29%	9.97

TABLE B7 | Optimization results for the F_1 family using RBF surrogates.

Iter.	Source	$\mathbf{t}_w$ [mm]	ε_{xy}	$\mathcal{D}_w$	$\mathcal{D}_f$	J
1	Surr.	[12.34, 14.34]	[0.04755, 0.06219]	[90.00, 90.00]	88.74	0.00
	FEM	[12.34, 14.34]	[0.04612, 0.06351]	[91.00, 94.60]	87.68	22.15
	Error	–	[3.12%, 2.07%]	[1.10%, 4.86%]	1.20%	22.15
2	Surr.	[12.68, 14.91]	[0.04531, 0.06034]	[90.00, 90.00]	87.53	0.00
	FEM	[12.68, 14.91]	[0.04499, 0.0598]	[88.74, 88.84]	86.81	2.42
	Error	–	[0.73%, 0.91%]	[1.42%, 1.31%]	0.83%	2.42
3	Surr.	[12.56, 14.79]	[0.04553, 0.06042]	[90.00, 90.00]	86.88	0.00
	FEM	[12.56, 14.79]	[0.04544, 0.06033]	[89.75, 89.84]	86.63	0.40514
	Error	–	[0.19%, 0.15%]	[0.28%, 0.18%]	0.29%	0.4051

TABLE B8 | Optimization results for the F_2 family using RBF surrogates.

Iter.	Source	t_w [mm]	ε_{xy}	$\mathcal{D}_w$	$\mathcal{D}_f$	J
1	Surr.	[15.50, 20.00]	[0.05243, 0.04593]	[105.09, 73.10]	96.41	507.38
	FEM	[15.50, 20.00]	[0.05211, 0.0454]	[105.82, 72.74]	98.50	880.86
	Error	–	[0.60%, 1.16%]	[0.69%, 0.50%]	2.12%	373.48
2	Surr.	[14.80, 18.93]	[0.05458, 0.04902]	[109.57, 78.18]	96.74	700.74
	FEM	[14.80, 18.93]	[0.05377, 0.04931]	[109.14, 78.71]	96.78	689.16
	Error	–	[1.52%, 0.59%]	[0.39%, 0.67%]	0.04%	11.58
3	Surr.	[14.77, 18.95]	[0.05395, 0.04908]	[109.61, 78.49]	96.64	688.53
	FEM	[14.77, 18.95]	[0.054, 0.04906]	[109.69, 78.51]	96.63	691.06
	Error	–	[0.10%, 0.05%]	[0.07%, 0.03%]	0.00%	2.53

TABLE B9 | Optimization results for the F_3 family using RBF surrogates.

Iter.	Source	t_w [mm]	ε_{xy}	$\mathcal{D}_w$	$\mathcal{D}_f$	J
1	Surr.	[5.94, 8.38, 9.28]	[0.04571, 0.05098, 0.05057]	[90.00, 90.00, 90.00]	71.81	0.00
	FEM	[5.94, 8.38, 9.28]	[0.04632, 0.04484, 0.05324]	[89.24, 82.01, 85.72]	69.82	13.04
	Error	–	[1.32%, 13.70%, 5.00%]	[0.86%, 9.75%, 5.00%]	2.85%	13.04
2	Surr.	[5.69, 7.97, 9.02]	[0.04677, 0.04855, 0.05551]	[90.00, 90.00, 90.00]	70.49	0.00
	FEM	[5.69, 7.97, 9.02]	[0.04928, 0.04787, 0.05344]	[94.17, 88.58, 88.69]	68.47	20.11
	Error	–	[5.09%, 1.42%, 3.86%]	[4.43%, 1.60%, 1.47%]	2.95%	20.11
3	Surr.	[5.81, 7.88, 8.98]	[0.04722, 0.04867, 0.05416]	[90.00, 90.00, 90.00]	69.17	0.00
	FEM	[5.81, 7.88, 8.98]	[0.04703, 0.04785, 0.05415]	[90.02, 89.07, 90.02]	68.85	0.93372
	Error	–	[0.42%, 1.70%, 0.01%]	[0.02%, 1.05%, 0.02%]	0.46%	0.9337

TABLE B10 | Optimization results for the F_4 family using RBF surrogates.

Iter.	Source	t_w [mm]	ε_{xy}	$\mathcal{D}_w$	$\mathcal{D}_f$	J
1	Surr.	[7.21, 9.27, 9.82]	[0.04351, 0.04034, 0.05378]	[90.00, 90.00, 90.00]	71.35	0.00
	FEM	[7.21, 9.27, 9.82]	[0.03892, 0.04601, 0.05525]	[82.20, 85.06, 90.33]	70.86	12.85
	Error	–	[11.80%, 12.31%, 2.66%]	[9.49%, 5.80%, 0.37%]	0.69%	12.85
2	Surr.	[6.81, 9.02, 9.65]	[0.0419, 0.04824, 0.05385]	[90.00, 90.00, 90.00]	69.63	0.00
	FEM	[6.81, 9.02, 9.65]	[0.04228, 0.04629, 0.05421]	[90.55, 89.06, 90.38]	69.09	1.39
	Error	–	[0.92%, 4.21%, 0.66%]	[0.61%, 1.05%, 0.42%]	0.78%	1.39

TABLE B11 | Optimization results for the F_5 family using RBF surrogates.

Iter.	Source	t_w [mm]	ε_{xy}	$\mathcal{D}_w$	$\mathcal{D}_f$	J
1	Surr.	[5.98, 7.29, 8.53, 6.73, 5.00]	[0.06013, 0.05438, 0.05386, 0.06869, 0.05284]	[90.00, 90.00, 90.00, 90.00, 74.98]	70.51	15.02
	FEM	[5.98, 7.29, 8.53, 6.73, 5.00]	[0.05632, 0.0526, 0.05986, 0.07083, 0.05884]	[84.59, 92.60, 89.02, 80.76, 73.89]	78.89	38.48
	Error	–	[6.77%, 3.39%, 10.02%, 3.02%, 10.19%]	[6.40%, 2.80%, 1.10%, 11.44%, 1.48%]	10.62%	23.46
2	Surr.	[5.71, 7.44, 8.51, 6.15, 5.00]	[0.05915, 0.05193, 0.06105, 0.07748, 0.05922]	[90.00, 90.00, 90.00, 90.00, 73.32]	80.01	16.68
	FEM	[5.71, 7.44, 8.51, 6.15, 5.00]	[0.05934, 0.05266, 0.05789, 0.07678, 0.05857]	[92.17, 89.57, 86.36, 90.05, 72.42]	77.82	26.36
	Error	–	[0.32%, 1.38%, 5.47%, 0.91%, 1.11%]	[2.36%, 0.47%, 4.21%, 0.06%, 1.25%]	2.81%	9.69
3	Surr.	[5.77, 7.38, 8.37, 6.18, 5.00]	[0.05805, 0.05275, 0.05965, 0.07628, 0.05832]	[90.00, 90.00, 90.00, 90.00, 72.26]	77.60	17.74
	FEM	[5.77, 7.38, 8.37, 6.18, 5.00]	[0.05914, 0.05314, 0.05986, 0.07616, 0.05681]	[89.36, 89.65, 89.64, 90.73, 72.63]	77.89	19.26
	Error	–	[1.84%, 0.73%, 0.34%, 0.14%, 2.66%]	[0.72%, 0.39%, 0.41%, 0.80%, 0.51%]	0.37%	1.51