A generalized autoregressive model for time-variant dynamical systems

Styfen Schär[1*], Chiara Nardin[1,2], Stefano Marelli[1] and Bruno Sudret[1]
[1] *Chair of Risk, Safety and Uncertainty Quantification,ETH Zurich, 8093 Zurich, Switzerland - {schaers,cnardin,marellis,sudretb}@ethz.ch*

[2] *Department of Civil, Environmental and Mechanical Engineering, University of Trento, 38132 Trento, Italy - {chiara.nardin}@unitn.it*

Abstract. Surrogate models are a staple tool for the reliability analysis of complex engineering systems, because they enable the use of otherwise computationally prohibitive algorithms, such as Monte-Carlo simulation. Their prevalence in the context of nonlinear dynamic engineering systems, however, is significantly lower, because approximating their time-dependent response is generally considered more challenging than for their static counterparts. A powerful class of surrogate models is provided by autoregressive with exogenous input models, but their use is mostly limited to system identification tasks, as they tend to underperform in pure prediction tasks. Recent advances in data-driven autoregressive modeling, have vastly improved their prediction capabilities for engineering applications, but they still rely on the assumption of time-invariant dynamics at the core of NARX. In this contribution we demonstrate how the mNARX$^+$ framework can efficiently handle time-variant dynamical systems, such as, e.g., civil or mechanical structures subject to plastic damage or fatigue accumulation, making it an ideal tool for their reliability assessment.

Keywords: Surrogate Modeling, Time-Variant Dynamical Systems, Autoregressive Modeling, mNARX, Uncertainty Quantification

1. Introduction

Structural reliability analysis has reached a high level of maturity, providing rigorous frameworks to assess the safety of engineering systems under uncertainty. When the underlying computational models are highly complex, direct simulation methods such as Monte Carlo Simulation (MCS) become computationally intractable, necessitating the use of fast-to-evaluate surrogate models (Moustapha et al., 2022).

For time-invariant problems, where the quantity of interest is a scalar or a low-dimensional vector, classical surrogate modeling techniques have become standard tools. Methods such as sparse polynomial chaos expansions (PCE) (Blatman and Sudret, 2011; Lüthen et al., 2021) and Kriging (Santner et al., 2003) allow for the accurate approximation of the input-to-output map with a limited computational budget. These techniques effectively replace the expensive solver within the reliability loop, enabling efficient estimation of failure probabilities, especially when paired with

© 2026 by the author(s). This work is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

active learning paradigms borrowed from the machine learning literature (Echard et al., 2011; Marelli and Sudret, 2018; Moustapha et al., 2022).

However, many critical engineering systems are inherently dynamical, *e.g.* offshore wind turbines, vehicle suspensions, or structures evolving under highly variable external loads, such as seismic excitations. In such cases, the model output is a time-dependent vector field or trajectory defined over a specific time domain. Rather than learning the global map directly, it can be significantly more effective to emulate the discrete-time evolution of the system, switching to a more state-space-like approach. This leads naturally to the class of Nonlinear AutoRegressive with eXogenous inputs (NARX) models (Billings, 2013).

By learning the recursive law governing the system, rather than the full trajectory, NARX models construct predictions based on the immediate past history of the input and the output, thereby naturally enforcing causality. Despite their utility, classical standard NARX formulations struggle heavily with the selection of optimal time lags and the processing of high-dimensional inputs, which often leads to instability over long time periods (Billings, 2013). Highly complex systems, such as wind turbines operating under turbulent winds or tall buildings subject to seismic loads, introduce additional challenges due to active control systems, uncertain structural parameters, etc.

To overcome these limitations, advanced surrogate models based on modified autoregressive approaches have recently been proposed, notably the manifold-NARX (mNARX) (Schär et al., 2024) and functional-NARX ($\mathcal{F}$-NARX) (Schär et al., 2025) frameworks. The mNARX approach constructs surrogates on a reduced-dimensionality manifold of auxiliary state variables, that can handle both high-dimensionality exogenous inputs, and highly nonlinear system responses. This is achieved by embedding expert knowledge into the regression formulation, by introducing a sequence of auxiliary NARX models of intermediate response quantities that are informative w.r.t. the desired quantity of interest (Schär et al., 2024). Conversely, $\mathcal{F}$-NARX mitigates the arbitrary nature of lag selection, long-memory effects and dependence on sampling rate, by extracting functional temporal features from continuous time windows (Schär et al., 2025). More recently, the mNARX$^+$ framework (Schär et al., 2026) was introduced to bridge these methods, automating the construction of the auxiliary manifold through data-driven feature selection, thereby removing the previous reliance on domain expertise.

In this paper, we capitalize on the mNARX methodology to handle the case of non-time-invariant engineering systems, by introducing auxiliary response variables that can track the state of the system over time. This allows us create an autoregressive manifold that can naturally account for time-varying states, e.g. due to progressive damage accumulation. The proposed approach ensures accurate trajectory predictions even after the structural integrity of the system has been partially compromised, an essential ingredient for efficiently quantifying their reliability and resilience.

2. Methodology

To model time-variant engineering systems undergoing progressive change, we need a robust mathematical framework that can capture and keep track of complex, nonlinear and time-dependent behaviors. This section outlines the theoretical foundations of autoregressive modeling, detailing the mNARX and $\mathcal{F}$-NARX frameworks before establishing the theoretical basis for our damage-tracking approach.

2.1. Standard NARX Framework

Consider a computational model of a dynamical system denoted by $\mathcal{M}$, that models some response quantity of interest $y(t)$. The system is subjected to a time-dependent, multi-dimensional exogenous excitation vector $\boldsymbol{x}(t)$. The model output over a time axis $\mathcal{T}$ is then given by:

$$y(t) = \mathcal{M}(\boldsymbol{x}(\mathcal{T} \leq t)). \tag{1}$$

Within this section, we also assume that the time axis $\mathcal{T} \in [0, T]$ is discretized regularly with a sampling interval $\delta t \ll T$.

A standard NARX model assumes that the current output $y(t)$ depends on its own past values (autoregressive terms) and the current and past values of the exogenous input. This discrete-time relationship is formally expressed as:

$$y(t) = \mathcal{H}\Big(y(t-\delta t), \ldots, y(t-n_y\delta t), \boldsymbol{x}(t), \boldsymbol{x}(t-\delta t), \ldots, \boldsymbol{x}(t-n_x\delta t)\Big) + e(t) \tag{2}$$

where n_x and n_y represent the maximum time lags for the exogenous inputs and system responses, respectively, $\mathcal{H}$ is a nonlinear mapping function, and $e(t)$ represents the residual error. In systems with high-dimensional exogenous inputs or those requiring long-term memory, the input space for $\mathcal{H}$ grows exponentially, exacerbating the curse of dimensionality and severely complicating surrogate training (Billings, 2013).

2.2. Functional Feature-based NARX ($\mathcal{F}$-NARX)

To overcome the limitations of classical discrete-time NARX models, the authors introduced the $\mathcal{F}$-NARX approach (Schär et al., 2025). It represents a major step away from discrete time steps by exploiting the functional nature of system trajectories over continuous time windows.

In a typical $\mathcal{F}$-NARX implementation, principal component analysis (PCA) is used to project historical time windows of the exogenous inputs and past outputs onto a set of orthogonal basis functions. The corresponding set of coefficients $\boldsymbol{\xi}(t)$ are then used instead of the input lags in an autoregressive setting (Schär et al., 2025).

In other words, Eq. (1) is rewritten as:

$$y(t) \approx \widehat{\mathcal{M}}(\boldsymbol{\xi}(\mathcal{T} \leq t)). \tag{3}$$

The model in Eq. (3) is then approximated with a standard NARX approach, as in Eq. (2).

This approach successfully maps continuous trajectory segments to interpretable time-dependent features, rather than isolated discrete points in time, fundamentally mitigating the numerical instabilities associated with very small time steps and long memory limits (Schär et al., 2025).

2.3. Manifold-NARX (mNARX)

To efficiently approximate the response of complex dynamical systems over extended time periods, the mNARX approach further constructs the surrogate on a problem-specific, reduced-order manifold of auxiliary state variables $\boldsymbol{\zeta}(t)$ (Schär et al., 2024).

By leveraging the underlying physics of the system as well as prior expert knowledge, mNARX decomposes the full problem into a hierarchical sequence of smaller sub-problems, which are then combined to form an input manifold $\boldsymbol{\zeta}(t)$ suitable for autoregressive modeling. More in detail, the manifold $\boldsymbol{\zeta}(t) = \{\tilde{\boldsymbol{x}}(t), \hat{\boldsymbol{z}}(t)\}$ is formed by two components: $\tilde{\boldsymbol{x}}(t)$, a set of features extracted from the input excitation $\boldsymbol{x}(t)$, e.g. through dimensionality reduction or spectral decomposition, and $\hat{\boldsymbol{z}}(t)$, a set of auxiliary quantities of interest that may be related to the output quantity of interest (QoI) $y(t)$, such as auxiliary quantities, or even quantities post-processed from the model response $y(\mathcal{T} < t)$. In general, the latter are not directly available from the input excitation, but need to be treated as additional model responses, and therefore modeled themselves as autoregressive surrogates:

$$\hat{z}_k(t) \approx \widehat{\mathcal{M}}_k \left(\tilde{\boldsymbol{x}}(\mathcal{T} \leq t), \hat{z}_1(\mathcal{T} < t), \ldots, \hat{z}_{k-1}(\mathcal{T} < t)\right). \tag{4}$$

Once the complete input manifold is available, the output QoI can be modeled autoregressively:

$$\hat{y}(t) = \widehat{\mathcal{M}}_y\Big(\boldsymbol{\zeta}(\mathcal{T} \leq t), \hat{y}(\mathcal{T} < t)\Big). \tag{5}$$

As demonstrated in Schär et al. (2024), this approach enables the efficient and accurate surrogate modeling of complex dynamical systems, even in the presence of high dimensional random inputs and active control systems. Further extensions to this algorithm include fully data-driven automatic construction of the input manifold (Schär et al., 2026), and long term functional output prediction, including estimation of epistemic error (Song et al., 2026).

2.4. Using mNARX to track damage in time-variant systems

In their native formulation (Eq. (2)), NARX models implicitly rely on the time-invariant nature of the underlying system. This is what allows them to accurately predict future states based on a finite model memory. In many real situations, however, dynamical systems vary over time, e.g. due to progressive damage accumulation, thus requiring autoregressive models to deal with potentially infinite memory. The same limitation also applies to the $\mathcal{F}$-NARX and mNARX models presented in the previous sections.

However, one of key features of the mNARX approach is that it can encode expert information about the system state directly in the input manifold $\boldsymbol{\zeta}(t)$, and use it to improve its prediction capability. If we can define one or more suitable *system-state trackers* $\boldsymbol{D}(t)$, that either represent or correlate to some indicator of the current state of the system, they can be included into the auxiliary manifold components $\hat{\boldsymbol{z}}(t)$ to keep track of the system state outside of its autoregressive window. Examples of such trackers include cumulative damage indices in structural systems, such as the Park–Ang damage index for earthquake-excited structures (Park and Ang, 1985); rainflow counting in fatigue (ASTM, 2017); effective stiffness degradation inferred from force-displacement hysteresis in nonlinear systems (Ibarra et al., 2005); or shifts in dominant response frequencies associated with stiffness loss (Lieven et al., 2001; Fassois and Sakellariou, 2006). Depending on the application, these indicators may be directly available as additional simulator outputs or may need to be estimated through post-processing of the system response

The resulting augmented manifold then becomes:

$$\tilde{\boldsymbol{\zeta}}(t) = \{\tilde{\boldsymbol{x}}(t), \hat{\boldsymbol{z}}^*(t)\}, \quad \hat{\boldsymbol{z}}^*(t) = \left\{\hat{\boldsymbol{z}}(t), \hat{\boldsymbol{D}}(t)\right\}, \tag{6}$$

where $\hat{\boldsymbol{D}}(t)$ is an autoregressive approximation of the chosen damage trackers $\boldsymbol{D}(t)$. This formulation ensures that the surrogate model explicitly parametrizes the state trackers $\boldsymbol{D}(t)$, continuously updating its internal representation of the system. In other words, the framework dynamically adapts, providing stable predictions across different system states, without the need to modify the overall formulation of the mNARX methodology.

3. Application: portal frame subject to high amplitude seismic excitation

3.1. Computational Model

The model adopted is based on, and slightly modified from, the reinforced concrete (RC) portal frame example provided in the OpenSeesPy manual (see the RCFrameEarthquake example in OpenSeesPy documentation (OpenSeesPy, 2024)).

More specifically, a one-bay, base-fixed RC portal frame with dimensions 144 in $\times$ 366 in and rectangular cross-sections (16 in $\times$ 25 in) was implemented as a two-dimensional nonlinear finite element model within the OpenSeesPy framework (OpenSeesPy Development Team, 2024). The columns were modeled using `forceBeamColumn` elements to capture both material and geometric nonlinearities, including P-Δ effects. A Gauss–Lobatto integration scheme with nine integration points per element was employed. The cross-sections were discretized into fiber components, consisting of: (i) confined concrete core ($f_c' = -6.0$ ksi); (ii) unconfined cover concrete ($f_c' = -5.0$ ksi); and (iii) longitudinal reinforcing steel ($f_y = 60$ ksi). Concrete behavior was represented using the `Concrete01` (Kent–Scott–Park) model, while reinforcing steel was modeled with the `Steel02` material with elastic modulus $E = 30000$ ksi, strain-hardening ratio $b = 0.02$, and standard transition parameters $R_0 = 18.0$, $c_{R1} = 0.925$, and $c_{R2} = 0.15$, as illustrated in Figure 1. Instead, the beam was modeled as a linear elastic `elasticBeamColumn` element. Constant axial loads of $P = 180$ kips were applied at both column tops through preliminary static analysis.

Differently from the reference example, the damping model was implemented using Rayleigh damping, calibrated to provide approximately 5% critical damping in the first two modes, following standard recommendations for RC structures, see (Chopra, 2007).

Finally, the dynamic response of the system is evaluated by imposing seismic excitation as a `UniformExcitation` load pattern, through transient analysis.

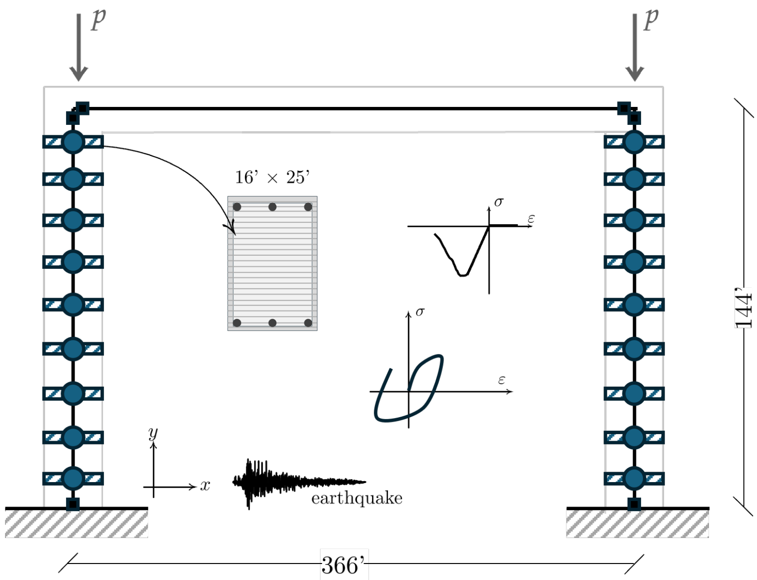

Figure 1. RC portal frame model and column fiber section, including Gauss–Lobatto integration (9 points) and material models for confined/unconfined concrete and steel.

3.2. Synthetic input ground motions

To assess the validity of the proposed framework, an ensemble of synthetically generated ground motions was employed as time-varying input. Specifically, synthetic accelerograms were generated using the site-based stochastic ground motion model (SGMM) proposed in (Su, 2025)[1]. This approach enables the generation of statistically compatible ground motions by fitting a stochastic model to a dataset of recorded motions representative of the seismic hazard of interest.

To construct the reference dataset, 50 real ground motions were selected following the methodology proposed by Baker and Lee (2018)[2], using the Pacific Earthquake Engineering Research Center (PEER) NGA-West2 database. The objective Baker's selection algorithm is to identify a subset of recorded ground motions whose response spectra optimally match a prescribed target distribution. This is achieved through an optimization procedure that minimizes the error between the spectra of the selected records and the target spectrum over a broad period range (0.1 s to 10 s), thereby ensuring both spectral compatibility and statistical representativeness. The selection was refined through a weighted optimization process ($w_{mean} = 1.0, w_{std} = 2.0, w_{skew} = 0.3$), prioritizing the match of the target standard deviation to better focus on capturing the inherent record-to-record variability.

The overall procedure consisted of a two-stage filtering approach. First, a target median response spectrum was defined based on a reference rupture scenario characterized by moment magnitude $M_w = 6.5$, Joyner-Boore distance $R_{jb} = 11$ km, and site condition $V_{s30} = 800$ m/s. Second, metadata-based constraints were imposed to ensure physical consistency of the selected motions.

[1] Available at the corresponding GitHub repository: `https://github.com/MaijiaSu/Toolbox_Site_based_GMMs`
[2] Available at the corresponding GitHub repository: `https://github.com/bakerjw/CS_Selection`

In particular, records were restricted to magnitudes between 6.0 and 8.2 and rupture distances between 10 and 50 km. This procedure ensures that the selected records are consistent with the seismological characteristics of the target scenario.

Once the dataset was defined, the SGMM was calibrated using the selected 50 records, allowing the joint probability distribution of the key ground motion parameters to be characterized, *i.e.* Arias intensity $\log(I_a)$, effective durations D_{0-5}, D_{5-30}, D_{30-45}, D_{45-75}, D_{75-95}, D_{95-100}, dominant frequency $\omega(t_{mid})$, rate of change of the filter frequency $\omega'(t_{mid})$, filter bandwidth $\zeta(t_{mid})$, and corner frequency fc. Subsequently, synthetic ground motion time histories were sampled from the resulting model, resulting in an ensemble of 10,000 synthetic records.

3.3. Problem setup

In this case study, we construct an mNARX surrogate model to predict the strain in the upper steel reinforcement bars at the base section (see Figure 1), denoted by $\varepsilon(t)$ and hereafter referred to as *steel top strain*. The input excitation is given by the ground motion acceleration introduced in Section 3.2, denoted by $x(t)$.

Based on the quantity of interest $\varepsilon(t)$, we define a damage indicator $D(t)$ that quantifies exceedance of a threshold strain value ε_{th}. Specifically, we consider a cumulative damage measure defined as

$$D(t) = \int_0^t \left(\varepsilon^+(\tau)\right)^2 \, d\tau, \tag{7}$$

where $\varepsilon^+(t)$ denotes the strain exceeding the threshold:

$$\varepsilon^+(t) = \max\left(|\varepsilon(t)| - \varepsilon_{\text{th}},\, 0\right). \tag{8}$$

The threshold $\varepsilon_{\text{th}} = 0.2\%$, below which strain excursions are assumed elastic and non-damaging, is inspired by the implicit yield threshold of Park and Ang (Park and Ang, 1985). The continuous-time integral formulation, which avoids cycle counting, follows the spirit of continuum damage mechanics (Lemaitre and Chaboche, 1990).

A visual illustration of this damage indicator is provided in Figure 2, which shows the steel top strain together with the corresponding accumulated damage. The latter can be interpreted as the squared area (highlighted in grey) outside the threshold bounds $\pm\varepsilon_{\text{th}}$.

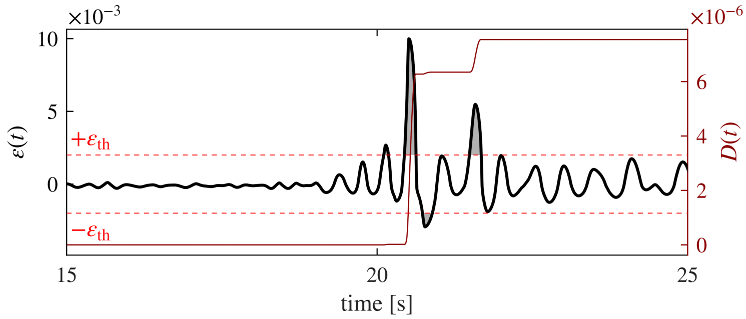

Figure 2. Example strain response and corresponding accumulated damage.

The experimental design consists of 100 simulated ground motion records and their corresponding strain responses and damage histories, which are used to train the surrogate model. For validation, an independent out-of-sample dataset comprising 10,000 ground motion records and corresponding strain responses is employed. Note that during validation the surrogate does not have access to the true underlying damage values.

3.4. Surrogate construction

Within the mNARX framework from Section 2.3, we employ an $\mathcal{F}$-NARX model as introduced in Section 2.2, using a multilayer perceptron (MLP) as the underlying function approximator for the mapping $\widehat{\mathcal{M}}$ in Eq. (3). The MLP consists of two hidden layers with 256 neurons each and ReLU activation functions.

The memory length of the $\mathcal{F}$-NARX model is set to 0.5 seconds (corresponding to $n = 50$ time steps). Temporal features are extracted using principal component analysis (PCA), following Schär et al., with a target explained variance of 90%. This results in 12 features for the exogenous input $\boldsymbol{\xi}_x(t)$ and 3 features for the autoregressive component $\boldsymbol{\xi}_\varepsilon(t)$.

In addition, the surrogate incorporates the cumulative damage $D(t)$ as defined in Section 3.3. During training, the $\mathcal{F}$-NARX model $\widehat{\mathcal{M}}$ learns the mapping

$$\varepsilon(t) \approx \widehat{\varepsilon}(t) = \widehat{\mathcal{M}}\big(\boldsymbol{\xi}_x(t), \boldsymbol{\xi}_\varepsilon(t), D(t - \delta t)\big), \tag{9}$$

with

$$\boldsymbol{\xi}_x(t) = \mathcal{K}_x\left(x(t), x(t - \delta t), \ldots, x(t - n\delta t)\right), \tag{10}$$

$$\boldsymbol{\xi}_\varepsilon(t) = \mathcal{K}_\varepsilon\left(\varepsilon(t - \delta t), \varepsilon(t - 2\delta t), \ldots, \varepsilon(t - n\delta t)\right), \tag{11}$$

where $\mathcal{K}_x(\cdot)$ and $\mathcal{K}_\varepsilon(\cdot)$ denote feature extraction operators applied to the exogenous input and autoregressive histories, respectively. In this work, both operators are implemented using PCA as mentioned above. Note that the true strain and damage values are provided during training.

In contrast, when predicting the system output for new excitations, the damage is iteratively calculated on the fly from Eq. (7), using the predicted strain $\hat{\varepsilon}(t)$:

$$\varepsilon(t) \approx \widehat{\varepsilon}(t) = \widehat{\mathcal{M}}\big(\boldsymbol{\xi}_x(t), \boldsymbol{\xi}_{\widehat{\varepsilon}}(t), \widehat{D}(t - \delta t)\big), \tag{12}$$

with

$$\boldsymbol{\xi}_{\widehat{\varepsilon}}(t) = \mathcal{K}_\varepsilon\left(\widehat{\varepsilon}(t - \delta t), \widehat{\varepsilon}(t - 2\delta t), \ldots, \widehat{\varepsilon}(t - n\delta t)\right), \tag{13}$$

and

$$\widehat{D}(t) = \sum_{k=0}^{\frac{t}{\delta t}-1} \left(\widehat{\varepsilon}^{+}(k\delta t)\right)^2 \delta t, \quad \widehat{\varepsilon}^{+}(t) = \max\big(|\widehat{\varepsilon}(t)| - \varepsilon_{\mathrm{th}}, 0\big). \tag{14}$$

To train the model in Eq. (9), we use an 80/20 train–test split of approximately 500,000 regression samples (corresponding to 100 simulations with an average length of 5,000 time steps). The model parameters are optimized using the Adam optimizer with a learning rate of 0.001 and a batch size of 512. Training is continued until convergence of the mean absolute error (MAE), which is assumed when the validation error does not decrease for 50 consecutive epochs. Based on this criterion, convergence is reached after approximately 400 epochs.

For comparison, we also train a baseline $\mathcal{F}$-NARX model with identical architecture but without the damage input. This allows us to highlight the limitations of finite-memory autoregressive models when applied to time-variant systems, and to demonstrate the benefit of mNARX with explicit damage tracking in such scenarios. Using the same training procedure as with the mNARX model, the baseline model converges after approximately 450 epochs.

3.5. Results

The overall performance of the mNARX model with damage tracking and baseline model without damage tracking on the out-of-sample validation dataset is shown in Figure 3.

Figure 3 (left) presents the distribution of the normalized mean squared error (NMSE) on a logarithmic scale. The NMSE for the i-th realization is defined as

$$\mathrm{NMSE}^{(i)} = \frac{1}{N_t^{(i)}\,\mathrm{Var}(\boldsymbol{\varepsilon}^{(i)})} \sum_{k=0}^{N_t^{(i)}-1} \left(\varepsilon^{(i)}(k\delta t) - \widehat{\varepsilon}^{(i)}(k\delta t)\right)^2, \tag{15}$$

where $N_t^{(i)}$ denotes the number of time steps and $\mathrm{Var}(\boldsymbol{\varepsilon}^{(i)})$ the variance of the time-series $\boldsymbol{\varepsilon}^{(i)}$.

Figure 3 (right) compares the predicted and true peak absolute strain values $\max(|\varepsilon(t)|)$ on logarithmic axes.

The surrogate model incorporating damage tracking significantly outperforms the baseline model without damage tracking. Most realizations exhibit NMSE values below 0.5, with a median of 0.07, whereas the baseline model yields substantially larger errors, with a median NMSE of 2.6 and many realizations exceeding 10.

This improvement is also evident in the peak strain predictions. The damage-aware surrogate closely follows the diagonal, indicating accurate reconstruction of extreme responses, while the baseline model shows a systematic deviation. Interestingly, the benefit of damage tracking is most pronounced for peak strains below approximately 5%, whereas both models exhibit comparable performance at higher strain levels.

We interpret this overall improvement at lower ground motion amplitudes as a direct consequence of the lack of damage tracking: the baseline surrogate effectively learns an "average" system response between undamaged and damaged systems, corresponding to an "effective" intermediate damage state. Consequently, it performs poorly when the true system is either undamaged or strongly damaged. On the contrary, the availability of a damage indicator in the input manifold allows the damage-aware surrogate to effectively separate the different system operational regimes, resulting in significantly improved model expressivity and accuracy.

The reduced accuracy at high strain levels for both models can be attributed to the limited representation of this regime in the training data: only three simulations exhibit peak strains above 5%, and only approximately 1.7% of all time steps in the experimental design show such high strains.

These trends are further illustrated by individual time histories shown in Figure 4. The top panels correspond to cases without accumulated damage, while the bottom panels show strongly

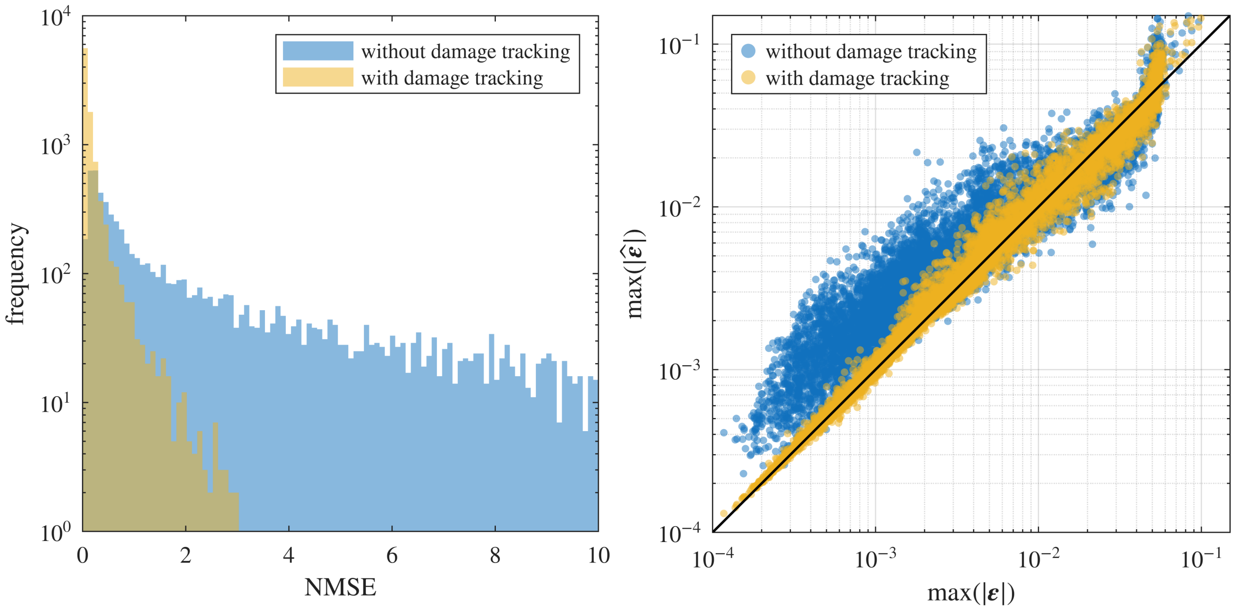

Figure 3. Distribution of the normalized mean squared error (NMSE) on the out-of-sample validation dataset (left), and comparison of predicted versus true peak absolute strain (right) for the surrogate model with and without damage tracking. The histogram (left) uses a logarithmic y-axis, while the scatter plot (right) uses logarithmic scaling on both axes.

damaged systems. The NMSE for these four representative simulations, for both the surrogate with and without damage tracking, are provided in Table I.

Table I. NMSE, as defined in Eq. (3.5), corresponding to the traces shown in Figure 4.

	Top left	Top right	Bottom left	Bottom right
With damage tracking	0.0066	0.0210	0.1444	0.3357
Without damage tracking	6.0327	1.9077	0.2937	0.9130

In the undamaged case, both models accurately follow the ground truth at early time instants; however, the baseline surrogate gradually deviates over time, whereas the damage-aware model remains accurate throughout the entire duration of the simulations. In the damaged case in the bottom right panel, the baseline model fails to represent the accumulation of residual strain, as it effectively approximates an averaged system response and therefore reverts toward zero-mean behavior. In contrast, the damage-tracking surrogate captures both the transient response and the evolving permanent deformation.

These observations confirm that incorporating damage tracking enables the surrogate to adapt to structural state changes, which leads to significantly improved predictive accuracy across a wide range of conditions.

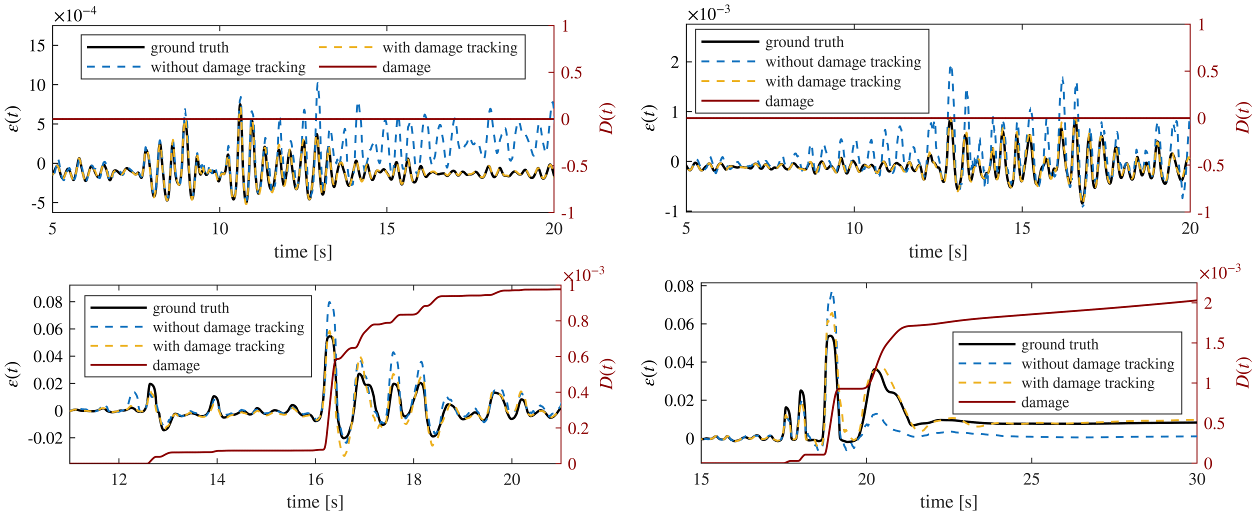

Figure 4. Example time histories. Top: simulations without accumulated damage. Bottom: simulations with significant damage. The ground truth is shown in black, the surrogate without damage tracking in blue, and the surrogate with damage tracking in yellow. The damage tracker $D(t)$ is shown in red.

3.6. Discussion and conclusion

In this paper, we introduced a generalized framework that enables the use of finite-memory autoregressive models to represent time-variant dynamical systems. Building upon the mNARX (Schär et al., 2024) and $\mathcal{F}$-NARX (Schär et al., 2025) methodologies, the proposed approach augments the autoregressive structure with auxiliary state-tracking variables that capture the evolving properties of the system.

The effectiveness of the method is demonstrated on a reinforced concrete portal frame subjected to severe, realistic ground motion excitations. In this context, we compared a baseline $\mathcal{F}$-NARX model without state tracking to an mNARX model incorporating a damage-based tracking variable. The results clearly show that the baseline model struggles to generalize across different structural states, performing poorly in both undamaged and highly damaged regimes. In contrast, the damage-aware surrogate is able to accurately capture the system response across the full range of conditions.

These findings indicate that augmenting a NARX-type model with appropriate state-tracking variables provides a practical and effective means of extending finite-memory autoregressive models to time-variant systems, which would otherwise require infinite memory.

Future work will focus on exploring more advanced and potentially multiple tracking variables, as well as automated strategies for their selection, leveraging already available information to extract other quantities of interest, such as displacements or forces, in order to improve robustness and generalization capabilities of the final surrogate.

Acknowledgements

This research is supported by the Marie-Sklodowska Curie program and the REACTIS project, GA no. 101147351. Views and opinions expressed are however those of the authors only and do not necessarily reflect those of the European Union, or REA or any sponsor. Neither the European Union nor the granting authority can be held responsible for them.

References

ASTM: 2017, 'Standard practices for cycle counting in fatigue analysis'. (E1049-85).

Baker, J. W. and C. Lee: 2018, 'An improved algorithm for selecting ground motions to match a conditional spectrum'. *Journal of Earthquake Engineering* **22**(4), 708–723.

Billings, S. A.: 2013, *Nonlinear system identification: NARMAX methods in the time, frequency, and spatio-temporal domains*. Chichester, West Sussex, United Kingdom: John Wiley & Sons, Inc.

Blatman, G. and B. Sudret: 2011, 'Adaptive sparse polynomial chaos expansion based on least angle regression'. *J. Comput. Phys.* **230**(6), 2345–2367.

Chopra, A. K.: 2007, *Dynamics of structures: theory and applications to earthquake engineering*. Pearson Prentice Hall, 3 edition.

Echard, B., N. Gayton, and M. Lemaire: 2011, 'AK-MCS: An active learning reliability method combining Kriging and Monte Carlo Simulation'. *Structural Safety* **33**(2), 145–154.

Fassois, S. D. and J. S. Sakellariou: 2006, 'Time-series methods for fault detection and identification in vibrating structures'. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **365**(1851), 411–448.

Ibarra, L. F., R. A. Medina, and H. Krawinkler: 2005, 'Hysteretic models that incorporate strength and stiffness deterioration'. *Earthquake Engineering & Structural Dynamics* **34**(12), 1489–1511.

Lemaitre, J. and J. Chaboche: 1990, *Mechanics of solid materials*. Cambridge University Press.

Lieven, N. A. J., D. J. Ewins, C. R. Farrar, S. W. Doebling, and D. A. Nix: 2001, 'Vibration–based structural damage identification'. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **359**(1778), 131–149.

Lüthen, N., S. Marelli, and B. Sudret: 2021, 'Sparse polynomial chaos expansions: Literature survey and benchmark'. *SIAM/ASA Journal on Uncertainty Quantification* **9**(2), 593–649.

Marelli, S. and B. Sudret: 2018, 'An active-learning algorithm that combines sparse polynomial chaos expansions and bootstrap for structural reliability analysis'. *Structural Safety* **75**, 67–74.

Moustapha, M., S. Marelli, and B. Sudret: 2022, 'Active learning for structural reliability: Survey, general framework and benchmark'. *Structural Safety* **96**, 102174.

OpenSeesPy: 2024, 'Reinforced concrete frame earthquake analysis example'. OpenSeesPy Documentation, accessed March 19, 2026.

OpenSeesPy Development Team: 2024, 'OpenSeesPy: Python interface for OpenSees'. `https://openseespydoc.readthedocs.io`. Accessed: 2026-03-19.

Park, Y. and A. Ang: 1985, 'Mechanistic seismic damage model for reinforced concrete'. *Journal of Structural Engineering* **111**(4), 722–739.

Santner, T. J., B. J. Williams, and W. I. Notz: 2003, *The design and analysis of computer experiments*. Springer, New York.

Schär, S., S. Marelli, and B. Sudret: 2024, 'Emulating the dynamics of complex systems using autoregressive models on manifolds (mNARX)'. *Mechanical Systems and Signal Processing* **208**, 110956.

Schär, S., S. Marelli, and B. Sudret: 2025, 'Surrogate modeling with functional nonlinear autoregressive models (F-NARX)'. *Reliability Engineering & System Safety* p. 111276.

Schär, S., S. Marelli, and B. Sudret: 2026, 'mNARX+: A surrogate model for complex dynamical systems using manifold-NARX and automatic feature selection'. *Computer Methods in Applied Mechanics and Engineering* **449**, 118550.
Schär, S., S. Marelli, and B. Sudret: 2026, 'mNARX+: A surrogate model for complex dynamical systems using manifold-NARX and automatic feature selection'. *Computer Methods in Applied Mechanics and Engineering* **449**, 118550.
Song, Z., M. A. Valdebenito, S. Schär, S. Marelli, B. Sudret, and M. G. R. Faes: 2026, 'Probabilistic function-on-function nonlinear autoregressive model for emulation and reliability analysis of dynamical systems'. arXiv preprint.

Su, M.: 2025, 'Stochastic ground motion models for uncertainty quantification in earthquake engineering'. Phd thesis, University of Trento, Trento, Italy. Repository: `https://github.com/MaijiaSu/Toolbox_Site_based_GMMs`.