Chapter 6

Adaptive Reduced Order Model Selection in Engineering Decision Support

Konstantinos Vlachas, Antonios Kamariotis, and Eleni Chatzi

Abstract Model-based approaches offer a structured, physics-informed framework to support engineering decisions related to the operation & maintenance of infrastructure assets. The core element of such frameworks are numerical models that distill complex systems into computationally tractable forms, while capturing essential dynamic behavior, variability and uncertainty. Within this context, Reduced Order Models (ROMs) facilitate real-time operations by reconciling accelerated model evaluations with accurate estimations. This balance can be achieved through techniques that embed physical insight into the ROM structure allowing for extrapolation and generalization capabilities. However, the utility of ROMs hinges on a delicate trade-off between computational savings and fidelity relative to Full Order Models (FOMs); a critical consideration when model outputs inform high-stakes decisions. In this work, we leverage real-time sensor data to propose an adaptive approach to dynamically select the appropriate model resolution, allowing the representation to evolve in step with the monitored system. First, a set of candidate ROMs is derived, each incorporating a generic parametrization to account for environmental and operational variability as well as for evolving phenomena, such as damage onset/progression or deterioration processes. Next, a sequential Bayesian inference framework is developed for quantification of the posterior uncertainty associated with each considered candidate ROM of a given resolution, thereby guiding the selection of the most suitable model representation. In turn, the generated ROMs are employed as predictive surrogates within the assembled inference framework, which assimilates monitoring information in a recursive manner. As a result, the selection among the generated models is posed as a problem of decision-making under uncertainty. The objective function is formulated to favor the resolution that maximizes the expected utility, defined as a trade-off between two competing criteria: i) precision in terms of recovering target quantities of interest and ii) computational efficiency measured by runtime. Subsequently, the selected resolution can be employed to support response prediction tasks and inform decision-making, while explicitly quantifying the associated uncertainty. The proposed integrated methodology enables adaptive model management by determining whether to accept the current model's prediction, switch to a different ROM instance or parameter realization, or trigger an on-the-fly retraining process based on real-time vibration measurements.

Keywords Reduced order models (ROMs) · Bayesian inference · Uncertainty quantification · Continuous data integration · Decision support

Introduction

A typical challenge in twining of structural infrastructure assets are their large dimensionality and the underlying complex dynamics that are characterized by stochasticity and uncertainty. Thus, for the derived virtual assets to deliver on their promise, they often rely on models that are both precise and computationally tractable [1]. Without loss of generality, in this work we focus on model-based techniques, and specifically Reduced Order Models (ROMs) [2]. These offer a structured and physics-based framework that provides better interpretability and more reliable extrapolation compared to purely data-driven

Konstantinos Vlachas · Eleni Chatzi
Department of Civil, Environmental and Geomatic Engineering, ETH Zurich, Stefano-Franscini-Platz-5, 8093, Zurich, Switzerland
e-mail: vlachas@ibk.baug.ethz.ch; chatzi@ibk.baug.ethz.ch

Antonios Kamariotis
Hitachi Energy AG, Baden, Switzerland
e-mail: antoniskam@hotmail.com

Kristin B. Zimmerman, PhD (eds.), *Model Validation and Uncertainty Quantification, 2026, Vol. 3*, Conference Proceedings of the Society for Experimental Mechanics Series, https://doi.org/10.1007/97887-438-1442-9_6

surrogates when tasked to support engineering decisions related to the operation & maintenance (O&M) of infrastructure assets. This has been demonstrated via utilizing ROMs as forward simulators for an array of tasks that necessitate repeated model evaluations, such as fault detection [3], input estimation [4], or identification under uncertainty [5].

In this context, the serviceability of ROMs lies in navigating the balance between computational cost and fidelity relative to recovering target Quantities of Interest (QoI). In this work we tackle this challenge by proposing an adaptive framework that integrates monitoring data continuously to dynamically select the optimal model for the downstream task at hand. Existing hybrid-twinning methodologies that attempt to address such adaptations in changing environments typically rely on passive data assimilation [6, 7], or problem-specific approaches that use suitable indicators to trigger retraining or refinement steps [8, 9], thus limiting their adaptability to continuously evolving conditions. In short, any model adaptation or updating happens in a static and one-off manner and the updated model's fidelity is not weighed against operational constraints, leading to potentially sub-optimal representations. In contrast, our work introduces a value-of-information-driven model selection [10], which occurs in a continuous and online manner, enabling the recursive identification of the task-optimized model to be utilized. Thus, the proposed approach enables the recursive selection of the lowest-cost model that can deliver predictions of the required fidelity for the system of interest in the context of interest, maximizing the overall benefit at any given time [11].

The proposed implementation initiates by deriving ROMs of different fidelity levels. As mentioned, ROMs are employed given their ability to facilitate real-time operations by reconciling accelerated model evaluations with accurate estimations. The developed ROMs incorporate a generic parametrization scheme that accounts for environmental variability and evolving conditions like structural damage [12]. Next, the candidate ROMs are incorporated into a sequential Bayesian inference framework, which quantifies the posterior uncertainty associated with each considered model of a given resolution, thereby guiding the selection of the optimal representation [13]. The assimilation of sensing data happens recursively and in a continuous manner leading to an adaptive framework able to follow along the evolving asset. The corresponding selection process aims to maximize the expected utility by weighing two competing factors: how precisely we can recover our quantities of interest, and how efficiently we can compute them to maintain real-time feasibility. The developed integrated methodology is illustrated in Fig. 1 and is employed to decide whether to accept the hybrid model's prediction, switch to a different generated instance or parameter realization, or trigger a retraining event, thus enabling an adaptive treatment.

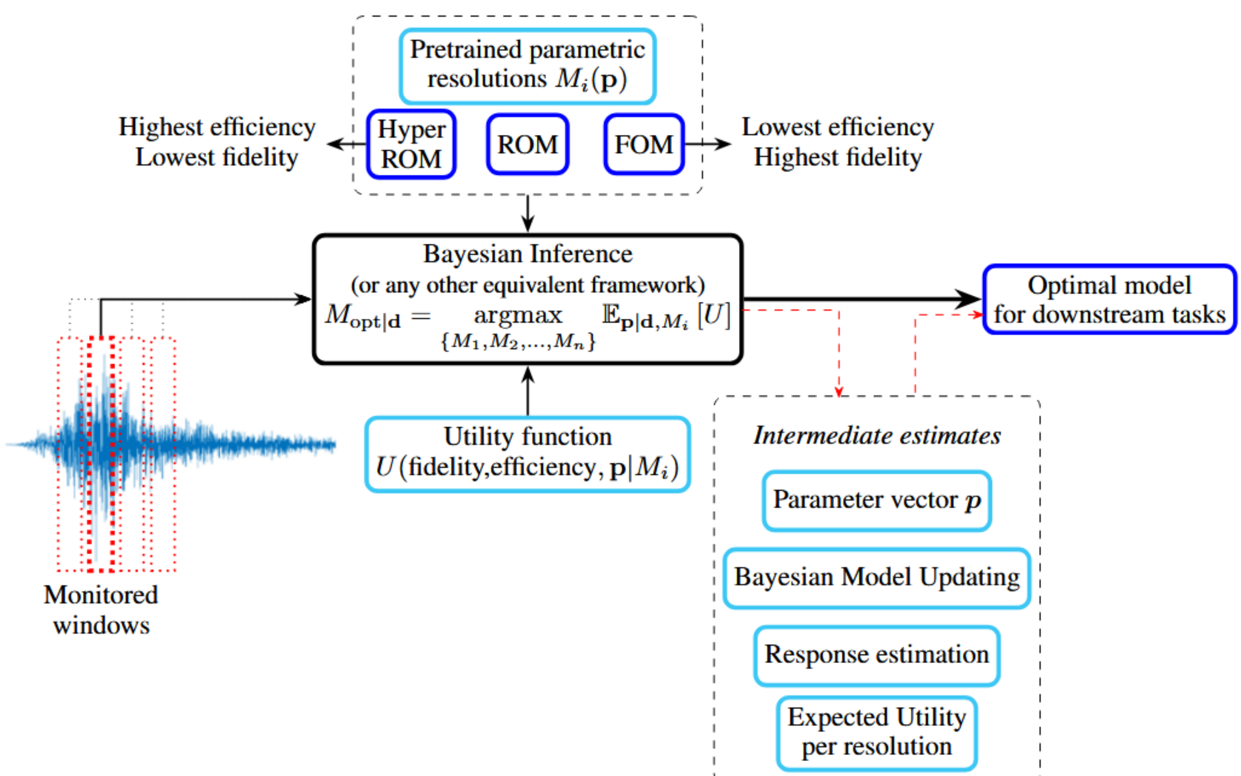

Fig. 1 Graphical abstract illustrating the proposed framework. The continuous monitoring and recursive model updating leads to an adaptive treatment

Problem Statement

Our work assumes availability of a Full-Order Model (FOM), which corresponds to a finite element (FE) model that is dependent on the input vector $\boldsymbol{p} = [p_1, ..., p_\mathrm{k}]^\mathrm{T} \in \Omega \subset \boldsymbol{R}^k$, which captures all system- and excitation-relevant parameters. Thus, the response is described by the following set of governing equations:

$$\boldsymbol{M}\ddot{\boldsymbol{u}}(t) + \boldsymbol{g}\left(\boldsymbol{u}(t), \dot{\boldsymbol{u}}(t), \boldsymbol{p}\right) = \boldsymbol{F}(t, \boldsymbol{p}), \tag{1}$$

where $\boldsymbol{u}(t) \in \boldsymbol{R}^n$ denotes the response in terms of displacements, $\boldsymbol{M} \in \boldsymbol{R}^{n\times n}$ the mass matrix, and $\boldsymbol{F}(t, \boldsymbol{p}) \in \boldsymbol{R}^n$ the induced excitation. For the sake of simplicity, the parametric dependency in the mass matrix has been dropped. Nonlinear behavior is captured via the restoring force term $\boldsymbol{g}\left(\boldsymbol{u}(t), \dot{\boldsymbol{u}}(t), \boldsymbol{p}\right) \in \boldsymbol{R}^n$. The full-order dimension n represents the total number of degrees of freedom and indirectly expresses the computational efficiency of the model, as necessary operations scale with n.

VpROM: A conditional Variational Autoencoder (cVAE) boosted ROM

The reduction approach employed herein relies on a Galerkin projection scheme and the underlying assumption that the dynamic behavior, namely the solution of Eq. 1, lies in a low-order subspace of dimension r, with $r \ll n$. Therefore, the following holds:

$$\boldsymbol{u}\left(\boldsymbol{p}\right) \approx \boldsymbol{V}\left(\boldsymbol{p}\right)\boldsymbol{q} \tag{2}$$

where $\boldsymbol{V} \in \mathbb{R}^{N\times r}$ represents the projection basis and $\boldsymbol{q} \in \mathbb{R}^r$ is the low-order coordinate vector. Via substitution of $\boldsymbol{u}$ into Eq. 1 and after multiplying the governing set with $\boldsymbol{u}^T$, thus performing a Galerkin projection, the following can be derived:

$$\tilde{\boldsymbol{M}}\ddot{\boldsymbol{q}}\left(t\right) + \tilde{\boldsymbol{g}}\left(\ddot{\boldsymbol{q}}, \dot{\boldsymbol{q}}, \boldsymbol{p}\right) = \tilde{\boldsymbol{F}}\left(\boldsymbol{p}, t\right) \tag{3}$$

where $\tilde{\boldsymbol{M}} = \boldsymbol{V}^T\boldsymbol{M}\boldsymbol{V}$, $\tilde{\boldsymbol{g}} = \boldsymbol{V}^T\boldsymbol{g}$ and $\tilde{\boldsymbol{F}} = \boldsymbol{V}^T\boldsymbol{F}$. A number of techniques have been proposed for the assembly of $\boldsymbol{V}$ [2]. Herein we employ the Proper Orthogonal Decomposition [12], which relies on a collection of FOM evaluations for a training set of parameters and assembles $\boldsymbol{V}$ via truncating the corresponding decomposition modes using suitable error measures. The Energy Conserving Mesh Sampling and Weighting hyper-reduction technique is also employed herein, which refers to a second-tier approximation required to tackle the computational bottleneck of updating $\tilde{\boldsymbol{g}}$ in Eq. 3 efficiently [14].

To address variability, a typical strategy is to assemble a collection of local projection bases $\boldsymbol{V}_i$, each corresponding to FOM snapshots for a specific realization $\boldsymbol{p}_i$, and utilize suitable interpolation or clustering techniques [15]. In this context, the proposed framework utilizes a nonlinear generative model in the form of a cVAE to capture a generalizable nonlinear mapping in the parameter-basis relation. This is shortly presented here and the reader can refer to [12, 16] for a more detailed discussion. First, a collection of local bases is obtained as described and the following system is solved in the least-squares sense:

$$\boldsymbol{V}_i\left(\boldsymbol{p}_i\right) = \boldsymbol{V}_{global} * \boldsymbol{X}_i\left(\boldsymbol{p}_i\right) \tag{4}$$

where $\boldsymbol{V}_i \in \mathbb{R}^{n\times r}$ is an instance of the collection, $\boldsymbol{V}_{global} \in \mathbb{R}^{n\times\tilde{r}}$ captures the dynamics across the entire domain and $\boldsymbol{X}_i \in \mathbb{R}^{\tilde{r}\times r}$ is a coefficient matrix. $\tilde{r}$ signifies the total number of truncated modes retained on $\boldsymbol{V}_{global}$. In turn, the proposed approach uses a conditional VAE as a nonlinear approximator of coefficients $\boldsymbol{X}$, and thus local projection-subspaces $\boldsymbol{V}$ and local ROMs, conditioned on known system parameters. Regarding the implementation aspects, the parametric features $\boldsymbol{p}$ are concatenated with the inputs of the VAE $\boldsymbol{X}$, and its latent space $\boldsymbol{Z}$. This is represented by an (inferred) probabilistic distribution, which enables the extraction of uncertainty bounds on the response estimates by sampling it and propagating the uncertainty on the estimates to the output through the ROM. After training the cVAE, a generative model is derived that can be sampled through its probabilistic latent space to produce the reduced basis coefficients $\boldsymbol{X}_i$ in Eq. 4. Concretely, once the trained decoder is available, and given the inferred parameter vector realization $\hat{\boldsymbol{p}}$ via Bayesian inference, the variational distribution of the latent space can be sampled. In turn, the generated samples are passed through the decoder portion and predictions of the coefficients $\boldsymbol{X}$ can be made. The generation process is illustrated in Fig. 2

Model selection via Bayesian decision theory

The framework outlined in the previous section can be employed to generate different model resolutions to assist with decision-making and virtual sensing tasks. As highlighted in Fig. 1, three candidates are considered here, each one corresponding to a trade-off between computational efficiency and fidelity: a) The full-order FE model of the system, termed

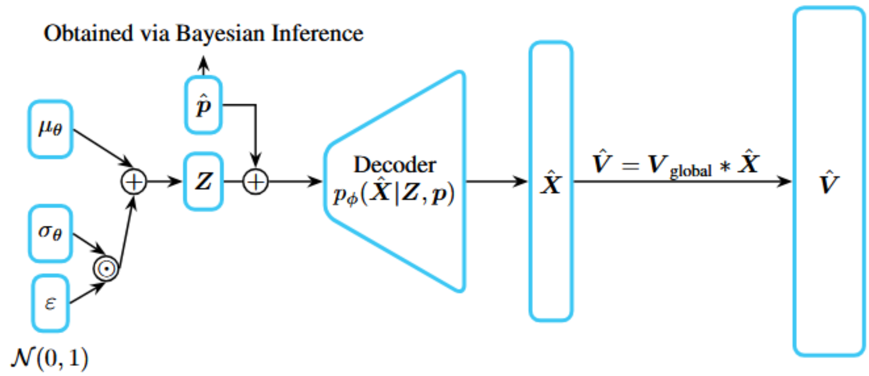

Fig. 2 The proposed framework in generation mode. Dimensionality serves only illustrative purposes

FOM, which offers the highest possible precision, b) the VpROM that strike a balance between precision and efficiency, and c) the H-VpROM, which corresponds to VpROM equipped with hyper-reduction, and potentially offers the highest efficiency. For selecting the best model available conditioned on incoming response measurements Bayesian decision theory is employed, whereby the optimal decision is to employ the representation that maximizes the expected utility [17]. The optimal model to be selected M_{opt} conditional on the monitoring data $\mathbf{d}$ can be obtained as:

$$M_{\text{opt}|\mathbf{d}} = \underset{\{M_1, M_2, \ldots, M_n\}}{\text{argmax}} \mathbb{E}_{\hat{\mathbf{p}}|\mathbf{d}, M_i}\left[U_i\right], \qquad (5)$$
$$\text{where } U_i = U(\text{fidelity}|M_i, \text{efficiency}|M_i, \hat{\mathbf{p}}|M_i)$$

where U is the considered utility function and $\hat{\mathbf{p}}$ the inferred model parameters. M_i represents the decision to employ the i-th candidate model for recovering the dynamic behavior of the engineered system of interest.

Two competing attributes that contribute to the utility U are considered here, namely the model's precision and efficiency. Since the Bayesian inference framework is also employed for parameter estimation, the utility is evaluated using readings from redundant sensor measurements that are not employed on the inference of $\hat{\mathbf{p}}$. The two branches of the utility function are illustrated in Fig. 3. In 3a underestimation is penalized more as it might lead to unexpected failures and catastrophic

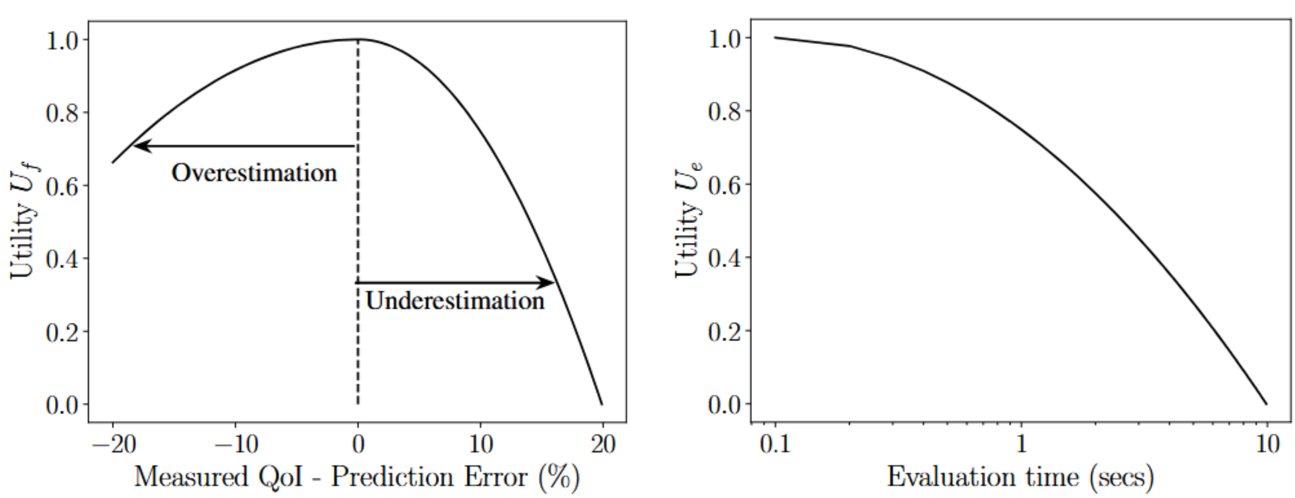

(a) Utility U_f as a function of the model fidelity

(b) Utility U_e as a function the model runtime

Fig. 3 Example utility functions for the considered competing attributes events. The overall utility can be computed as a weighted combination of the competing attributes as follows

$$U_i(\text{fidelity}|M_i, \text{efficiency}|M_i, \hat{\mathbf{p}}|M_i) = \gamma * U_i(\text{fidelity}|M_i, \hat{\mathbf{p}}|M_i) + (1-\gamma) * U_i(\text{efficiency}|M_i, \hat{\mathbf{p}}|M_i) \tag{6}$$

$$U_i(\text{fidelity}|M_i, \hat{\mathbf{p}}|M_i) = U_f(\hat{\mathbf{p}}|M_i) = \begin{cases} 1 - w_1 * \left(\frac{y_f - \hat{y}_f}{y_f}\right)^2 * (w_1 * \epsilon_{max})^{-1}, & (y_f - \hat{y_f}) > 0 \\ 1 - w_2 * \left(\frac{y_f - \hat{y}_f}{y_f}\right)^2 * (w_2 * \epsilon_{max})^{-1}, & (y_f - \hat{y_f}) \leq 0 \end{cases} \tag{7}$$

$$U_i(\text{efficiency}|M_i, \hat{\mathbf{p}}|M_i) = U_e(\hat{\mathbf{p}}|M_i) = 1 - \frac{\log^2(t_{M_i})}{\log^2(t_{max})} \tag{8}$$

where y_f denotes the observed QoI and $\hat{y}$ the model-approximated ones using the inferred $\hat{\boldsymbol{p}}$ and the selected resolution M_i. The variable t denotes the evaluation time. The weights w_1, w_2 are chosen properly to account for over- and under-estimation effects ($w_1 > w_2$) and to normalize the corresponding utility. Regarding the model updating task, the following probabilistic model is assumed for the discrepancy between the observed acceleration signals $\mathbf{y}$ and the model-predicted response $\hat{\mathbf{y}}$:

$$nRMSE(\hat{\mathbf{p}}, M_i) = \sqrt{\frac{\sum_{j=1}^{n_f}\left(\mathbf{y}_j - \hat{\mathbf{y}}_j\left(\hat{\mathbf{p}}, M_i\right)\right)^2}{\sum_{j=1}^{n_f}\mathbf{y}_j^2}} \in \mathbb{R}^{n_f} \sim \mathcal{N}\big(\mathbf{0}, \mathbf{\Sigma} = \text{diag}(c^2||\mathbf{y}||^2)\big) \tag{9}$$

where $\mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$ denotes the multivariate normal distribution with a zero-mean vector and covariance matrix $\mathbf{\Sigma}$. A diagonal covariance matrix is assumed, with the variance of each component assumed proportional to the L_2-norm of $\mathbf{y}$. The factor $c = 0.10$ can be regarded as a coefficient of variation, and its chosen value reflects the total prediction error. The likelihood function can be then written as:

$$L(\boldsymbol{p}; \mathbf{y}) \sim \mathcal{N}\big(\boldsymbol{\eta}; \mathbf{0}, \mathbf{\Sigma}\big), \tag{10}$$

where $\mathcal{N}(\cdot\,; \mathbf{0}, \mathbf{\Sigma})$ denotes the value of the MVN density function at a specified location. Thus, the expected fidelity utility U_f to employ model M_i can be approximated using n_s posterior samples as:

$$\mathbb{E}_{\hat{\mathbf{p}}|\boldsymbol{y}, M_i}[U_f(\hat{\mathbf{p}}|M_i)] \approx \frac{1}{n_s}\sum_{j=1}^{n_s} U_f\big(M_i, \hat{\mathbf{p}}_j\big), \tag{11}$$

The equivalent applies for U_e. Regarding the numerical implementation of the Bayesian inference components outlined in this section, the improved Transitional Markov Chain Monte Carlo (iTMCMC) method is employed [18].

Benchmark Shear Frame with Hysteretic Links

The numerical validation focuses on an open-source benchmark of a three-dimensional shear frame with nodal couplings that exhibit Bouc-Wen hysteresis, available in [19]. This example is chosen due to its relevance to real–life structures and its inherent ability to parametrically model multiple simultaneously activated instances of nonlinearity, representing damage or condition deterioration. The proposed framework is utilized to identify the optimal model that maximizes the cost-to-efficiency tradeoff, defined as the utility, to be used for response inference at any given time. The setup is summarized in Tab. 1 and a graphical illustration is offered in Fig. 4a, where the colored columns indicate the regions where evolving damage is introduced.

Due to the limited availability of experimental or actual monitoring measurements, simulated measurements are utilized instead. These are obtained via an independent numerical FE model, termed perturbed FOM, and by applying the following measures to avoid the so-called inverse crime:

- A finer mesh is used for the FE model representing the perturbed FOM
- A random perturbation of the Young's Modulus E is introduced.
- Measurement noise equal to 8% of the root-mean-square value of each acceleration signal is added.

Our setup assumes damage occurs in the basement or second floor columns, as illustrated in 4a, via modifying the stiffness in the corresponding Bouc-Wen springs. We also consider the Young modulus E as a varying parameter, to address

Table 1 Details of the case study and its numerical setup

Scenario	Structure behaves linearly at first. Structure and column joints get damaged (unevenly) due to ground motion.
Objective	Select model to be used for O&M downstream tasks
Monitoring Data	Nodal accelerations
	Four nodes monitored (blue crosses in Fig. 4a)
QoIs for Utility computation	Fidelity: Inter-story drift (Normalized RMSE)
	Efficiency: Avg. model evaluation time (s)
Utility functions	Equations 6-8, 3, $t_{max} = 10s, w_1 = 3, w_2 = 1, \epsilon_{max} = 20\%$
Monitoring window	1 second rolling windows

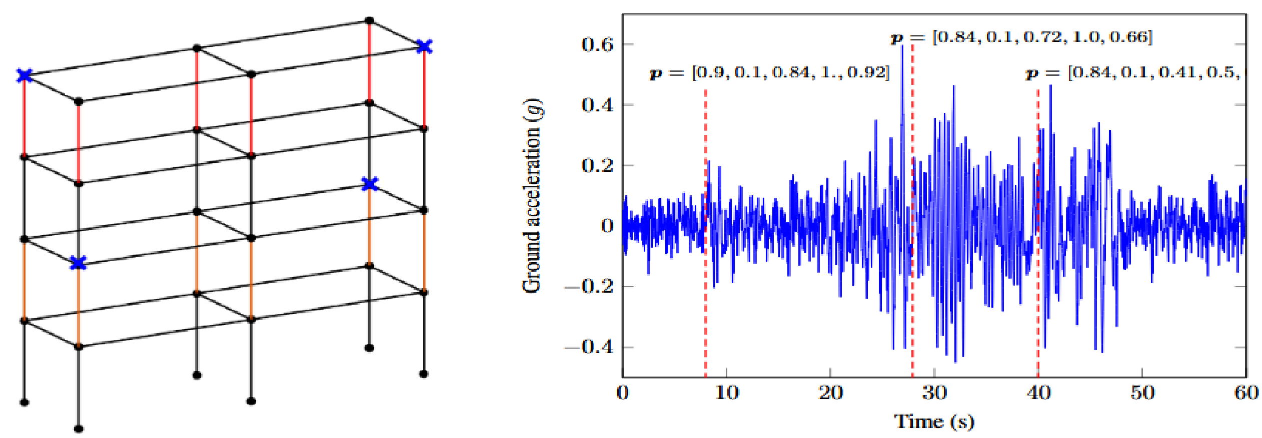

(a) Illustration of the shear frame (b) Testing input ground motion accelerogram

Fig. 4 Shear frame diagram and the testing ground motion accelerogram with damage events indicated by red lines. Initially the frame is assumed undamaged, and thus $\boldsymbol{p} = [1.0, 1.0, 1.0, 1.0, 1.0]$

global damage. Thus, the parameter vector to be inferred in Eq. 5 is $\boldsymbol{p} = [E, \alpha, k_1, k_2, k_3]$, where k_i indicates the k parameter in the links of floor i. Parameter α models the initially linear frame ($\alpha = 1.0$) that experiences damage via the activation of the nonlinear springs ($\alpha = 0.1$). The range and distribution of the parameters is summarized in Tab. 2. Regarding material properties, steel HEA cross sections are used, and the frame has $l = 15$ m, $w = 5$ m, and all stories have $h = 3.2$ m.

Forcing is applied to the frame as a base excitation scenario representing an earthquake. Base accelerograms are applied at an angle of $\theta = \pi/4$ with respect to the longitudinal axis. During training, synthetic accelerograms obtained from white noise signals are used, along with 500 training samples of $\boldsymbol{p}$. The framework is tested using the example input sequence in Fig. 4b and three condition deterioration events. Specifically, the frame is initially assumed healthy, and thus linear. After the first damage event, the hysteretic links are activated ($\alpha = 0.10$) with the corresponding stiffness coefficients ($k_1 = 0.84, k_2 = 1.0, k_3 = 0.92$). In the second event, the stiffness coefficients of the hysteretic springs in the first and last floor reduce unevenly, reflecting localized damage while the E modulus also reduces. Lastly, the third event reduces all three stiffness coefficients k_1, k_2, k_3 to values outside the operational regime that was considered for deriving the ROMs. This extreme case is designed to challenge the accuracy of the available reduced representations and test whether the framework will signal a retraining step.

Table 2 System properties and range of parameter values for the training of the ROMs. Variable α changes during operation to signify the activation of the nonlinear links, implying initiation of damage

Parameter:	E (×210 GPa)	$k_1, k_2 (\times 1e^8)$	α	β	γ	δ_η, δ_ν	A, w
Range:	[0.70, 1.0]	[0.6, 1.0]	$\in \{0.1, 1\}$	3.0	4.0	0.0	1.0

The total utility readings for the example scenario are reported in Fig. 5. The damage events are also illustrated using black dotted lines in the corresponding timestamps. As expected, the hyper-reduced model resolution, termed H-VpROM provides high accuracy response estimates under parametric variability and accelerated model evaluations, thus achieving the maximum utility before the $40s$ timestamp, despite the few low-amplitude oscillations observed. However, the FOM outperforms the ROM variants after the $40s$ damage event, when the integrated parametric traits assume values well-beyond the training domain. Despite the well-documented extrapolation capabilities of physics-based ROMs, this performance is expected, especially under extreme events. As a result, the proposed framework signals successfully that the FOM is to be preferred for tasks requiring high precision. In addition, a threshold can be defined in the utility curves that can serve as an indicator that signals the need for retraining a representation and triggers the corresponding calibration event. Naturally the optimal model selection that depends on the maximization of the utility function also reflects a decision on the importance of the competing attributes. Thus, in Fig. 5a, fidelity and precision are equally weighted, and the framework indicates that the H-VpROM and the FOM are the best models available before and after the $40s$ extreme event. For the sake of demonstration, in Fig. 5b computational runtime is favored over precision by setting $\gamma = 0.3$ in Eq. 6, and, thus, the VpROM seems the most suitable model after the $40s$ event.

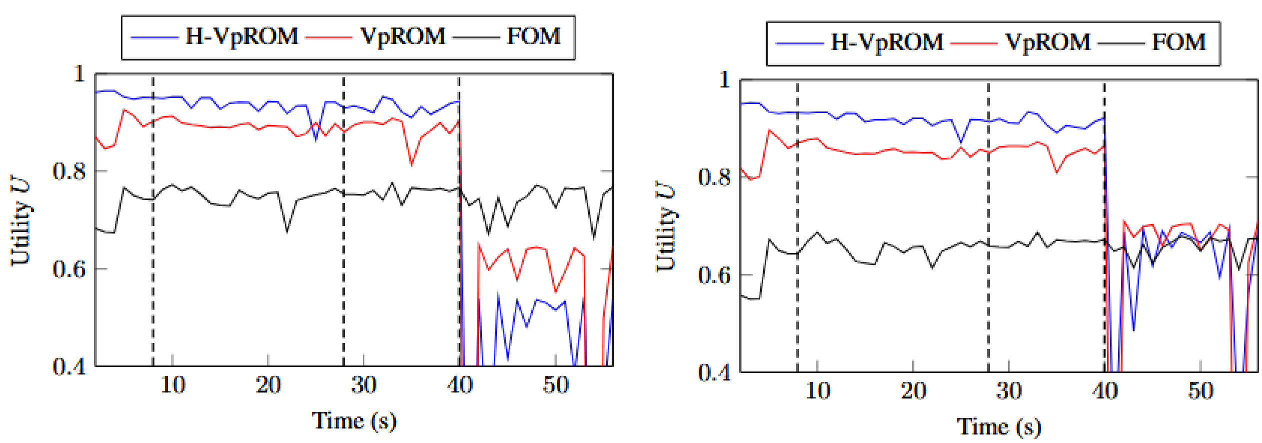

(a) Equally important utility attributes ($\gamma = 0.5$)

(b) Efficiency prioritized over precision ($\gamma = 0.3$)

Fig. 5 Total utility evaluations for different weighted combinations of the considered attributes in Eq. 6

To further scrutinize the performance of the proposed framework, the prior and posterior distributions of the k_1 uncertain model parameter are reported in Fig. 6a after the second damage event. As shown in 6a, the performance on the parameter

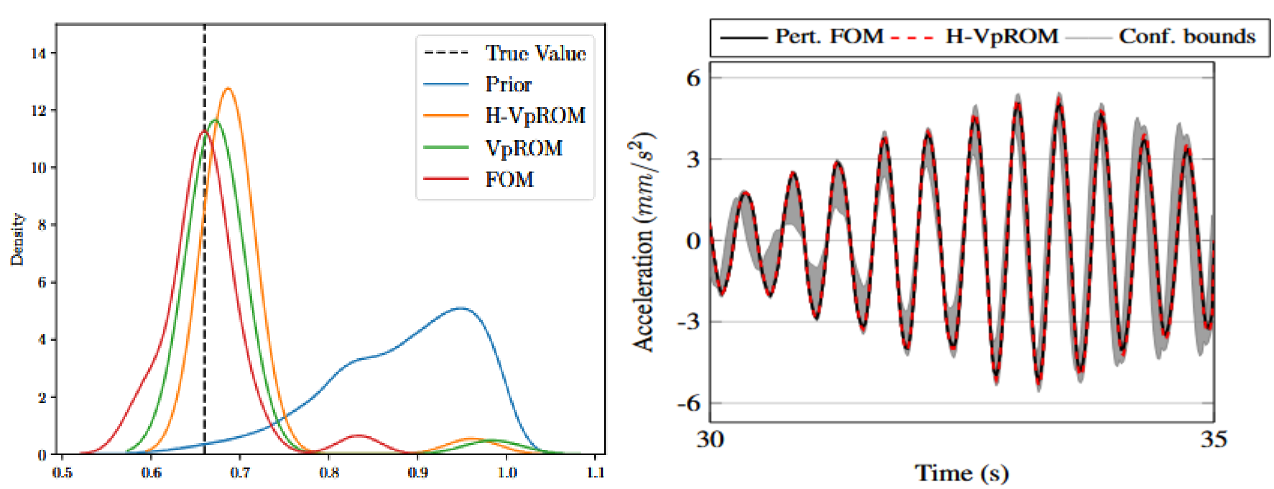

(a) Stiffness coefficient k_1 inference

(b) Response approximation using H-VpROM

Fig. 6 Output response approximation and error bounds at the 30s mark,including a 5s-ahead response prediction inference task depends on the underlying model functioning as forward simulator. The FOM exhibits the highest precision, followed by the VpROM and lastly the hyper-reduced H-VpROM. However, the H-VpROM is the optimal candidate, as reported in Fig. 5a, which further highlights why the utility should be the critical quantity for identifying the optimal model, as it weights both model fidelity and efficiency. The corresponding time history response approximation is reported in Fig. 6b, along with the uncertainty bounds obtained. Interestingly, these bounds contain the model's response and can be treated as indicative of the error and confidence that may be attributed to the response estimates, thus offering a quantification of the uncertainty that propagates into related decision tasks.

Discussion and Limitations

Our work addresses a key challenge in engineering decision support, namely the adaptive selection of the optimal model for downstream O&M tasks in an online and continuous manner. The proposed framework relies on a generative framework that outputs rapid but reliable ROMs of different resolutions conditioned on parametric inputs. These models are integrated as forward simulators of the system's dynamics into a Bayesian framework tasked to infer the underlying parametric traits from monitoring data, and, quantify the resulting utility of the models for the decision-maker. This is achieved via maximizing a decision-theoretic objective function that considers two competing attributes: i) precision, which takes into account the over/under-estimation of target quantities, and (ii) computational runtime.

Furthermore, the developed integrated methodology can be employed to suggest whether to accept a candidate model's prediction, switch to a different generated instance or parameter realization, or trigger a retraining event, thus enabling an adaptive treatment. The framework will be validated on large-scale examples and case studies involving experimental data as next to pinpoint its potential and limitations for real-time O&M.

Acknowledgments The authors gratefully acknowledge the funding from the European Commission under the Horizon Europe funding guarantee, for the projects 'ReCharged - Climate-aware Resilience for Sustainable Critical and interdependent Infrastructure Systems enhanced by emerging Digital Technologies' (grant agreement No: 101086413) and 'TURING - Trustworthy Unified Robust Intelligent Generative Systems' (grant agreement No: 101215032).

References

1. Chinesta, F. and Cueto, E. . Empowering engineering with data, machine learning and artificial intelligence: a short introductive review. *Advanced Modeling and Simulation in Engineering Sciences*, 9(1):21, 2022.
2. Benner, P. , Ohlberger, M. , Cohen, A. , and Willcox, K. . *Model reduction and approximation: theory and algorithms*. SIAM, 2017.
3. Azzi, M. J. and Farhat, C. . Enhanced Multimodal Nonparametric Probabilistic Method for Model-Form Uncertainty Quantification and Digital Twinning. *AIAA Journal*, pages 1–15, 2024.
4. Tatsis, K. E. , Agathos, K. , Chatzi, E. , and Dertimanis, V. K. . A hierarchical output-only Bayesian approach for online vibration-based crack detection using parametric reduced-order models. *Mechanical Systems and Signal Processing*, 167:108558, 2022.
5. Kamariotis, A. , Vlachas, K. , Ntertimanis, V. , Koune, I. , Cicirello, A. , and Chatzi, E. . On the consistent classification and treatment of uncertainties in structural health monitoring applications. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering*, 11(1):011108, 2025.
6. Moya, B. , Badías, A. , Alfaro, I. , Chinesta, F. , and Cueto, E. . Digital twins that learn and correct themselves. *International Journal for Numerical Methods in Engineering*, 123(13):3034–3044, 2022.
7. Rocha, I. B. , Meer, F. , van der, and Sluys, L. J. . An adaptive domain-based POD/ECM hyper-reduced modeling framework without offline training. *Computer Methods in Applied Mechanics and Engineering*, 358:112650, 2020.
8. Etter, P. A. and Carlberg, K. T. . Online adaptive basis refinement and compression for reduced-order models via vector-space sieving. *Computer Methods in Applied Mechanics and Engineering*, 364:112931, 2020.
9. Peherstorfer, B. and Willcox, K. . Online adaptive model reduction for nonlinear systems via low-rank updates. *SIAM Journal on Scientific Computing*, 37(4):A2123–A2150, 2015.
10. Kamariotis, A. , Chatzi, E. , and Straub, D. . Value of information from vibration-based structural health monitoring extracted via Bayesian model updating. *Mechanical Systems and Signal Processing*, 166:108465, 2022. Publisher: Elsevier.
11. Worden, K. , Cross, E. J. , Barthorpe, R. J. , Wagg, D. J. , and Gardner, P. . On digital twins, mirrors, and virtualizations: Frameworks for model verification and validation. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering*, 6(3):030902, 2020.
12. Vlachas, K. , Tatsis, K. , Agathos, K. , Brink, A. R. , and Chatzi, E. . A local basis approximation approach for nonlinear parametric model order reduction. *Journal of Sound and Vibration*, 502:116055, 2021.
13. Kamariotis, A. and Chatzi, E. . Bayesian decision-theoretic model selection for monitored systems. *arXiv preprint arXiv:2310.10485*, 2023.

14. Farhat, C. , Avery, P. , Chapman, T. , and Cortial, J. . Dimensional reduction of nonlinear finite element dynamic models with finite rotations and energy-based mesh sampling and weighting for computational efficiency. *International Journal for Numerical Methods in Engineering*, 98(9):625–662, 2014.
15. Amsallem, D. and Haasdonk, B. . PEBL-ROM: Projection-error based local reduced-order models. *Advanced Modeling and Simulation in Engineering Sciences*, 3:1–25, 2016.
16. Vlachas, K. , Simpson, T. , Garland, A. , Quinn, D. D. , Farhat, C. , and Chatzi, E. . Reduced order modeling conditioned on monitored features for response and error bounds estimation in engineered systems. *Mechanical Systems and Signal Processing*, 226:112261, 2025.
17. Keeney, R. L. and Raiffa, H. . *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge university press, 1993.
18. Betz, W. , Papaioannou, I. , and Straub, D. . Transitional markov chain monte carlo: observations and improvements. *Journal of Engineering Mechanics*, 142(5):04016016, 2016.
19. Vlachas, K. , Agathos, K. , Tatsis, K. E. , Brink, A. R. , and Chatzi, E. . Two-story frame with Bouc-Wen hysteretic links as a multi-degree of freedom nonlinear response simulator. In *5th Workshop on Nonlinear System Identification Benchmarks*, 2021.