

Learning large-scale industrial physics simulations

Fabien Casenave

In an industrial group like Safran, numerical simulations of physical phenomena are integral to most design processes. At Safran's corporate research center, we enhance these processes by developing fast and reliable surrogate models for various physics. We focus here on two technologies developed in recent years. The first is a physical reduced-order modeling method for non-linear structural mechanics and thermal analysis, used for calculating the lifespan of high-pressure turbine blades and performing heat analysis of high-pressure compressors. The second technology involves learning physics simulations with non-parameterized geometrical variability using classical machine learning tools, such as Gaussian process regression. Finally, we present our contributions to the open-source and open-data community.

1 Introduction

Consider an operator F that represents a complex numerical problem or an expensive experiment. Industrial design processes often require frequent evaluations of such operators in many-query tasks, such as parametric exploration, optimization, uncertainty quantification, or calibration. These tasks are intractable without modifications or simplifications, making them unsuitable for real-life industrial applications. For instance, optimizing the shape of an

aircraft wing to maximize finesse under a minimum thickness constraint involves many iterations of running a complex fluid solver, which is computationally intensive, see Figure 1.

A common approach to manage these tasks is by replacing the expensive operator F with a surrogate f . This surrogate is created by learning the behavior of F from data. A training database is constructed by selecting inputs using a design of experiments (DoE) algorithm and computing corresponding outputs by evaluating F . The surrogate, formed as a regressor trained on these input/output pairs, must be accurate enough to replace F in many-query tasks and fast enough to be practical for industrial applications.

For tabular problems where inputs and outputs are low-dimensional vectors ($\mathbb{R}^d \ni x \mapsto y = F(x) \in \mathbb{R}^p$), traditional regression methods like linear or polynomial regression, nearest neighbors, random forests, support vector regression, or Gaussian process regression can be used. This work focuses on learning physics problems, which involve complex, high-dimensional heterogeneous objects like meshes, boundary conditions, fields, or material models, unlike tabular problems. The inputs are the varying components of the physical setting in the many-query task, while the outputs are the predicted components used for decision-making.

For example, the AirFRANS physics learning problem [1] involves a mesh of the fluid domain, boundary conditions, the Reynolds-averaged Navier–Stokes equations, and the $k - \omega$ turbulence model.

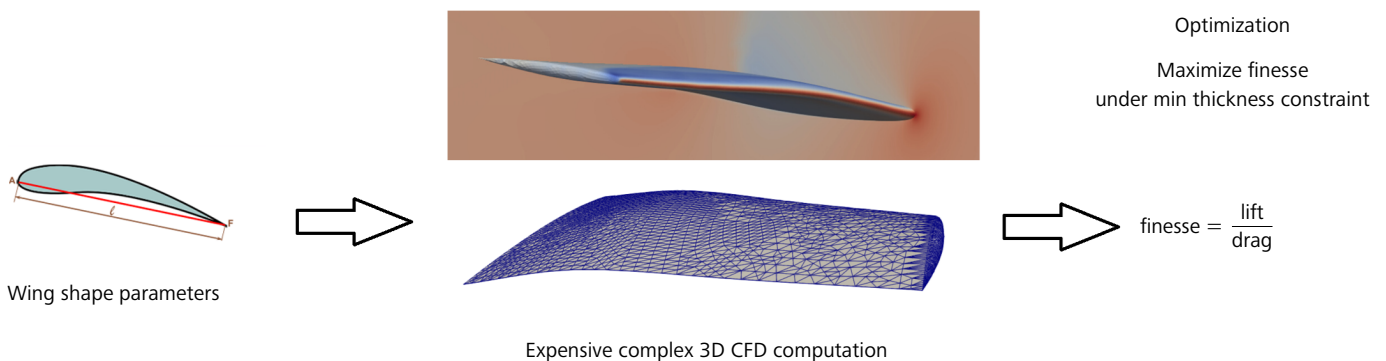


Figure 1. Optimization of the shape of an aircraft wing.

Here, the inputs include the mesh, angle of attack, and inlet velocity, while the outputs are the velocity, pressure, and turbulent velocity fields. The chosen turbulence model being fixed in the training dataset, it is not part of the inputs of the learning problem.

The advantage of physics learning problems over tabular problems with scalar inputs and outputs is the flexibility to change shapes beyond the training parameterization and the ability to modify the post-treatments of predicted solution fields after learning.

This article first discusses physical reduced-order modeling technologies developed at Safran in Section 2. Then, it introduces a machine learning method for non-parameterized geometrical variability in Section 3. Finally, it presents Safran's initiatives in the open-source and open-data communities in Section 4.

The article is a shortened version of the paper submitted to the Proceedings of the 9th European Congress of Mathematics, based on the work presented during the Felix Klein Prize Lecture on July 16th, 2024.

2 Physical reduced-order modeling

Physical reduced-order modeling (ROM) is a surrogate modeling technique that uses the underlying physics equations during the exploitation stage.

2.1 Linear data compression

During the training stage, we generate inputs using DoE algorithms and compute corresponding outputs by evaluating the operator F . We then apply a linear dimensionality reduction algorithm to these outputs to construct a reduced-order basis. In the exploitation stage, we use the same algorithms as in the high-dimension problem but restrict the solution to the subspace spanned by the reduced-order basis. For example, if F involves solving a high-dimensional initial boundary value problem, we replace the finite element basis with the reduced-order basis to keep using the Galerkin method, expressing the reduced solution as a linear combination of the reduced-order basis elements. Various methods have been proposed to adapt physical ROMs to different problems and simulation methods involved in F evaluations. One such method is the Reduced Basis method [20], which starts with a single input/output pair and enriches the data by selecting configurations where the current model makes the largest error, using fast and accurate error bounds [5] for accuracy assessments and to trigger enrichment steps.

The speed-up in physical ROMs comes from the reduced-order basis having a much smaller cardinality than the finite element basis, leading to smaller linear systems to solve in the Newton's algorithm during the exploitation stage. However, constructing these small systems can be costly, especially with non-linear problems requiring

numerical integration over the mesh. To achieve practical speed-up, precomputations and approximations (called hyperreduction) are performed during the training stage to reduce computational effort in the exploitation stage, see [4, Section 3.2] for more details, and [6, Section 1] for references on hyperreduction methods.

We focus on non-linear structural mechanics and non-linear transient heat problems, utilizing snapshot proper orthogonal decomposition (POD) for dimensionality reduction, which is parallelizable with domain decomposition. Additionally, we use a localized empirical cubature method (ECM), which derives reduced quadrature formulae independently for each subdomain. This involves solving an optimization problem to minimize the number of non-zero elements of the quadrature weights vector, using a greedy approach called non-negative orthogonal matching pursuit. This ensures the reduced operators retain the spectral properties of F , maintaining symmetry and positive definiteness of the tangent operator in the Newton's algorithm.

For non-linear structural mechanics and transient heat problems, we have demonstrated significant speed-ups, see respectively [4] and [6]. For instance, our physical ROM reduced the computation time for the lifetime prediction of 3D elastoviscoplastic high-pressure turbine blades from weeks to hours, and accelerated transient heat analysis of high-pressure compressors with complex boundary conditions, achieving errors compatible with industrial requirements.

In structural mechanics applications, quantities of interest often include dual quantities, such as plastic strain or von Mises stress. To reconstruct dual quantities over the entire mesh, the gappy POD method is used. By enforcing the reduced quadrature scheme to include well-chosen quadrature points, the gappy POD reconstruction can be made well-posed, see [4, Proposition 1]. Our numerical experiments revealed a strong correlation between the prediction error on dual quantities and the ROM-gappy-POD residual, which is defined as the error, on the reduced quadrature points, between the online computation of the corresponding quantity by the constitutive law solver and the gappy POD reconstruction. Thus, we introduced [3] an error indicator, which is a regressor trained to predict the error from the ROM-gappy-POD residual. This indicator successfully maintains the error below a chosen threshold by triggering enrichment steps.

Linear dimensionality reduction enables the application of the Galerkin method during the exploitation stage. However, this restriction can sometimes result in a reduced-order basis with a cardinality too large to achieve significant speed-up, even with hyperreduction. These situations, referred to as poorly reducible problems, will be addressed in the following section.

2.2 Piecewise linear data compression

The reducibility of the approximation of F by f linear dimensionality reduction is evaluated by the rate of decay of the Kolmogorov




Figure 3. Local ROM recommendation.

N -width of the output solutions with respect to the dimension N of an optimal approximation vector space. The Kolmogorov N -width measures the worst approximation error of a set of outputs by the best N -dimensional subspace. If the decay rate is fast, linear data compression methods are effective. If not, the reduced-order basis becomes too large, making the reduced problems slow to solve, indicating a poorly reducible problem. This challenge is addressed in Thomas Daniel's Ph.D. work [11].

To tackle poorly reducible problems, we focus on creating a dictionary of local physical ROMs using piecewise linear data compression, by clustering the outputs fields, see Figure 2. Consider




Figure 2. Piecewise linear data compression

the problem of predicting the lifetime of 3D elastoviscoplastic high-pressure turbine blades, now with the temperature loading field as input. To use a dictionary of local ROMs, we need to know in advance in which cluster of the outputs the solution is located, for us to use the adequate local ROM. To achieve this, we train a classifier to select the correct local ROM based on the 3D input temperature field, similar to how classifiers in computer graphics

recognize objects in images, see Figure 3. Such workflows are proposed in [13]. Unlike image classifiers, our situation involves few high-dimensional inputs. To improve the classifier training, we developed specialized feature selection and data augmentation techniques. This approach ensures accurate local ROM selection and effective handling of complex, high-dimensional input data, see [14].

Partitioning outputs using the L^2 -norm in the ambient solution space has been explored in the literature, but this does not guarantee low-dimensional approximation subspaces. In [12, Property 4.13], we demonstrate that the partitions of outputs minimizing the k -medoid cost function with sine dissimilarity are exactly the minimizers of the sum of a local variant of the Kolmogorov 1-width, weighted by the partition element's volume. This method, which considers the relative angles between outputs, offers a practical algorithm for optimally partitioning outputs for our purposes.

Consider a 2D advection problem where Gaussian processes with small (0.1) and large (1) amplitudes and various vertical positions ξ_2^0 move from left to right, see Figure 4. We illustrate the




Figure 4. Some output fields for the 2D advection problem [12].

clustering of these outputs using multidimensional scaling (MDS) with five clusters, comparing L^2 -norm and sine dissimilarity. With the L^2 -norm, all outputs with the small amplitude (0.1) are tightly packed at the center, likely grouping into the same cluster. This cluster contains all the independent directions of the outputs, hence




Figure 5. Number of POD modes for each local basis with respect to the number of clusters, for various accuracy criteria of the POD truncation, applied to the 2D advection problem [12].

has a Kolmogorov N -width not smaller than the one of the set of all outputs. Conversely, with sine dissimilarity, three connected components emerge, corresponding to the three considered ξ_2^0 values. At each position are located two outputs varying only by magnitude, indicating zero sine dissimilarity for outputs differing only in magnitude.

In Figure 5, we compare the effectiveness of clustering outputs using the L^2 -norm versus sine dissimilarity in constructing local ROMs of small dimensions. When using the L^2 -norm, one cluster ends up with the same dimension as a global ROM, specifically the cluster containing all small magnitude outputs. This indicates that no matter how many clusters are chosen, one cluster remains as poorly reducible as the entire set of outputs. Conversely, when using sine dissimilarity, the dimensions of all local ROMs decrease as the number of clusters increases. This demonstrates that clustering outputs using sine dissimilarity is beneficial, as it successfully reduces the ROM dimensions. Therefore, the choice of dissimilarity in output clustering is crucial, and the sine dissimilarity proves to be effective in this context.

In [15], a dictionary of local ROMs is developed to quantify the uncertainty of the plastic strain and von Mises stress in high-pressure turbine blades with respect to input temperature loading. First, we parameterize the variability of the input temperature field and construct two DoEs: one MaxProj of size 80 and one Sobol' of size 120. We evaluate the operator F for these 200 configurations, taking 7 days and 9 hours on 48 computer cores. Using the sine dissimilarity with the PAM k -medoid clustering algorithm, we partition the MaxProj DoE outputs into two clusters in 5 minutes. Next, we construct the two corresponding local ROMs using the POD-ECM technique presented in Section 2.1, which takes 5 hours on 24 computer cores. We label the 120 Sobol' DoE outputs by computing the sine dissimilarity with the two identified medoids in 5 minutes. Our feature selection algorithm takes 16 minutes on 280 computer cores, followed by training a logistic regression classifier with elastic-net regularization in 1 minute. This process results in a ROM dictionary of two local ROMs and a classifier ca-

pable of recommending the appropriate local ROM based on input temperature loading, making our surrogate f ready for use. Finally, we conduct probability density estimation of the plastic strain and von Mises stress in areas of interest on the turbine blade using 1,008 Monte Carlo evaluations of f . This yields relative errors of 1–2% and a speed-up factor of 600.

2.3 Limitations

Despite the successful application of physical ROMs to many equations and configurations, two significant challenges persist. First, non-parameterized geometrical variations make it difficult to compare outputs supported on different meshes. Second, global non-linear data compression prevents the use of the Galerkin method during the exploitation stage. Some authors have developed solutions to address these limitations while still constructing a surrogate f that involves solving the physics equations. However, these approaches often involve complex algorithms that can hinder practical speed-ups.

In the following section, we propose a surrogate f that incorporates non-linear data compression and manages non-parameterized geometrical variations, but does not rely on solving the physics equations during the exploitation stage.

3 Learning physical problems with non-parameterized geometrical variability

3.1 Mesh morphing Gaussian process

In our recent work [9], we introduce the mesh morphing Gaussian process (MMGP) technique, which integrates four key components: (i) mesh morphing, (ii) finite element interpolation, (iii) dimensionality reduction, and (iv) Gaussian process regression. This approach is designed to learn solutions to partial differential equations (PDEs) that involve geometric variations not explicitly parameterized.

Figure 6 illustrates the MMGP inference process for predicting a field of interest. The figure should be read from left to right: the left side shows various geometries corresponding to different samples, while the right side displays the target fields to be predicted on these geometries. For non-parameterized meshes, we perform a shape embedding process to convert meshes into learnable objects, by considering the coordinates of the mesh vertices as continuous fields. The left column of Figure 6 depicts the continuous field of the x -component of the coordinates, showing vertical iso-values. We then apply a deterministic morphing process to transform input meshes into a common shape, such as the unit disk shown in Figure 6. Each sample is converted into a mesh of the unit disk, modifying the coordinate fields according to the input mesh shapes. We then select a common mesh of the unit disk and project the coordinate fields from each morphed mesh onto this common mesh using finite element interpolation. This process

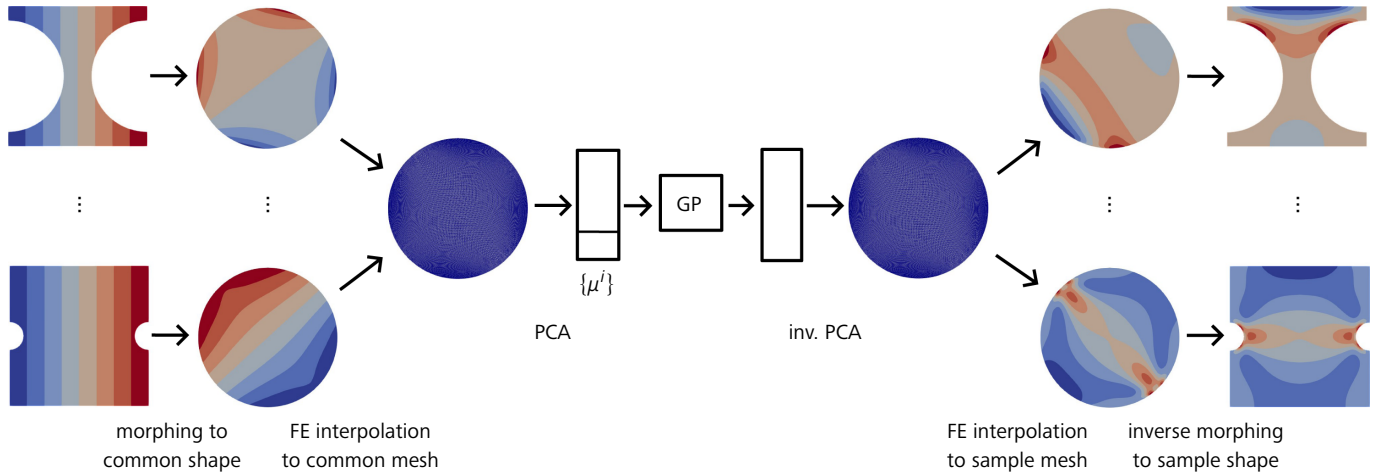


Figure 6. Illustration of the MMGP inference workflow for the prediction of an output field of interest [9].

results in all coordinate fields being represented on the same mesh, allowing us to use classical dimensionality reduction techniques, like principal component analysis (PCA), to generate low-dimensional vectors for the meshes. If scalar inputs are part of the learning problem, these are appended at the end of these vectors.

Similarly, a field embedding process is needed to convert variable-sized output fields into learnable objects. We apply the same transformations used for coordinate fields: morphing onto a common shape, finite interpolation onto a common mesh, and PCA compression. This process results in low-dimensional vectors representing our output fields of interest.

These deterministic pre-treatments effectively reduce the physics learning problem to a low-dimensional input/output tabular regression problem. The morphing and related steps act as the dimensionality reduction phase, which, while possibly highly non-linear, simplifies the machine learning task by converting large, variable-dimensional objects into smaller manageable ones. The deterministic nature of these pre-treatments is essential, as it simplifies the machine learning stage and avoids the complexity of dealing with large and variable-dimensional data directly.

We use Gaussian process regression for its high accuracy and ability to estimate predictive uncertainties effectively. For further details and numerical experiments demonstrating the efficacy of MMGP, including comparisons with recent deep learning technologies like MeshGraphNet and U-net graph convolutional networks, we refer to [9].

3.2 ML4PhySim competition

We participated in a competition [16] organized by IRT SystemX and partners, where we employed our MMGP technology. The challenge involved predicting pressure, velocity, and turbulent viscosity fields on the AirFRANS dataset [1], introduced in Section 1.

To achieve effective morphing, we first extended all 2D meshes to a common boundary box to handle irregular external boundaries, as shown in Figure 7.

We select the geometry of the first pretreated mesh from the training set as the common shape. We specify the morphing of chosen control points — on the intrados and extrados of the airfoil, the wake line, and the external boundary — and use a Radial Basis Function (RBF) morphing algorithm to compute the transformation of all other points. We project all fields, including input coordinates and output fields, onto the common mesh (the pretreated and morphed mesh from the first training sample). For dimensionality reduction, we used snapshot proper orthogonal decomposition (POD) as detailed in [9, Annex C]. The snapshot POD approach ensures that objects reconstructed from reduced representations preserve zero linear relationships, which is critical for maintaining zero boundary conditions, such as those for the velocity field at the airfoil surface. This helps to ensure that the zero boundary condition is respected in the predicted fields. We train separate Gaussian processes for each generalized coordinate associated with the snapshot POD of each output field. The Gaussian process kernel combines a constant term with an RBF kernel and a white noise term.

Some results are illustrated in Figure 8.




Figure 7. Illustration of the horizontal velocity for two samples of the AirFRANS dataset: the external boundaries are very irregular.




Figure 8. Illustration of the result on a test sample.

The competition score, ranging from 0 to 100, evaluates both the speed and accuracy of the surrogate model, including scalar quantities derived from the output fields (such as lift and drag coefficients) and their Spearman’s rank correlation coefficients, tested on two test sets, one of them being out-of-distribution. Table 1 shows that our method won first place, outperforming competitors who used advanced deep learning techniques, despite our approach using simpler and more classical tools.

3.3 Outlooks

MMGP is an effective method for approximating field and scalar outputs in complex physics problems with non-parameterized geometrical variability. It combines mesh morphing with finite element interpolation and uses shape embedding by reducing the dimensions of coordinate fields, simplifying the machine learning task. This approach enables efficient Gaussian process regression on reduced-dimensional data.

The method can handle large meshes, can be trained on CPU, is interpretable, and provides accurate predictions and predictive uncertainties. It is especially efficient for industrial part design, where configurations often have low intrinsic dimensions.

However, MMGP has some limitations: the morphing process must be customized for each use case, it assumes a fixed mesh topology, and speed-up is limited by the need for morphing and interpolation during inference. Future work will focus on addressing these issues by developing a generic morphing method, creating efficient approximate morphing procedures for inference [19], and deriving an optimal morphing strategy to minimize dimensionality further.

4 Open source and open data

Since the inauguration of the Safran corporate research center in 2015, we have implemented processes to make our research

rank	method	physics (30%)				ML-related (40%)					OOD generalization (30%)				global score (%)							
		physical criteria				accuracy (75%)					OOD accuracy (42%)		OOD physics (33%)									
		C_D	C_L	ρ_D	ρ_L	\bar{u}_x	\bar{u}_y	\bar{p}	\bar{v}_t	\bar{p}_s	speed-up (25%)	\bar{u}_x	\bar{u}_y	\bar{p}		\bar{v}_t	\bar{p}_s	C_D	C_L	ρ_D	ρ_L	speed-up (25%)
1	MMGP	●	●	●	●	●	●	●	●	●	27.4	●	●	●	●	●	●	●	●	●	28.08	81.29
2	GNN-FC	●	●	●	●	●	●	●	●	●	570.77	●	●	●	●	●	●	●	●	●	572.3	66.81
3	MINR	●	●	●	●	●	●	●	●	●	518.58	●	●	●	●	●	●	●	●	●	519.21	58.37
4	Bi-Trans	●	●	●	●	●	●	●	●	●	552.97	●	●	●	●	●	●	●	●	●	556.46	51.24
5	NeurEco	●	●	●	●	●	●	●	●	●	44.93	●	●	●	●	●	●	●	●	●	44.78	50.72

Table 1. ML4PhySim competition final leaderboard [17].

codes open-source and non-confidential datasets available as open-data. Our commitment to accessibility, quality, and robustness involves using DevOps tools and practices, such as Git for version control, code reviews, automated testing, deployment pipelines, and extensive documentation. We opt for permissive licenses like MIT or BSD3 for our code and CC-BY-SA for datasets. Our codes are hosted on GitLab, documentation on Read the Docs, and packages are distributed via conda-forge, when possible.

4.1 Open-source codes

One recent development is the PLAID (Physics Learning AI Data-model) library, which formalizes physics learning problems and manages complex datasets. It leverages the CGNS standard [10] for data handling, promoting consistency and ease of data sharing among practitioners.

The codes developed for the methods presented in the previous sections are open-source. Notably, Muscat [2, 29] is a li-

brary for mesh processing in finite element computations. PLAID [30] is used as a foundation for our physics learning developments. Physical ROM technologies are supported by two libraries: Mordicus [18], co-developed with the MOR_DICUS consortium, and genericROM [26], developed by Safran. The implementation of genericROM is non-intrusive, parallel with distributed memory, and can handle non-parameterized variability, see [25] for details on these notions and a description of Mordicus and genericROM. Additionally, MMGP, discussed in Section 3, is also available as open-source [28]. For tabular learning problems, Safran’s open-source platform Lagun [27] supports a range of tasks including design of experiments, surrogate modeling, sensitivity analysis, and optimization.

4.2 Open-data datasets

We are in the process of making our non-confidential datasets available as open-data in the PLAID format, facilitating their use

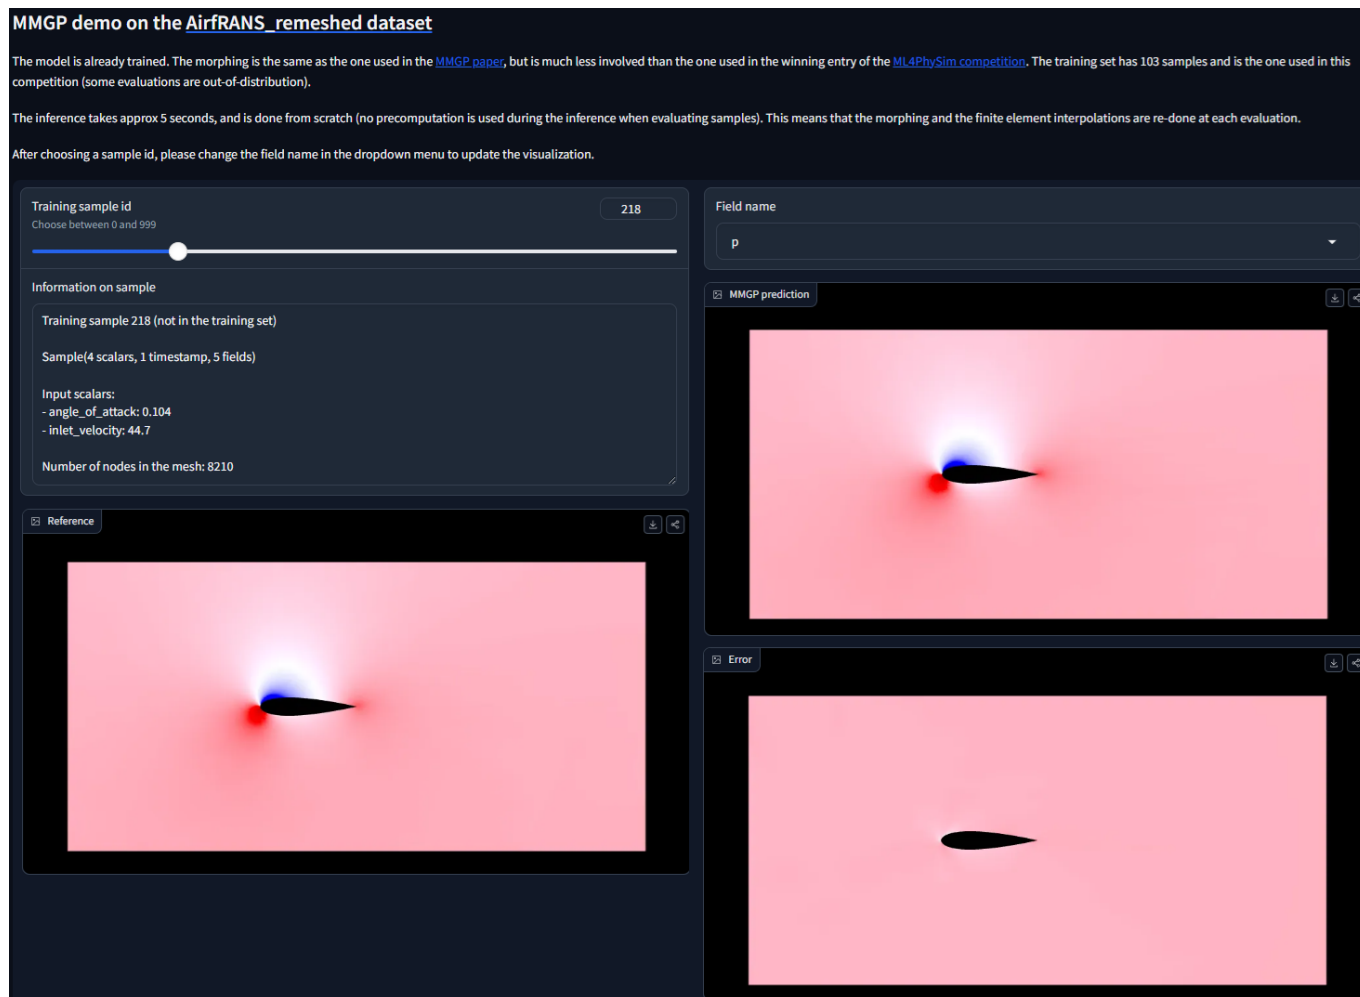


Figure 9. Hugging Face Spaces hosting a demo of our method MMGP applied on the physics learning problem of the ML4PhySim competition.

by the community. These datasets are shared through Zenodo and Hugging Face. Specifically, the three datasets used in our recent comparison of MMGP with advanced deep learning methods in [9] are accessible on Zenodo [7, 21, 22] and Hugging Face [8, 23, 24].

Additionally, we have distributed a demo of our MMGP method on Hugging Face’s Spaces platform, which was applied to the ML4PhySim competition problem presented in Section 3.2, see Figure 9. In this demo, users can select a sample ID, which triggers the MMGP surrogate model, which involves: morphing, finite element interpolation, projection onto snapshot POD modes, inference using a Gaussian process regressor, and finally, inverse snapshot POD and finite element interpolation. Predicted fields, reference data, and errors are plotted. Note that for real-time performance, the demo uses a remeshed dataset with meshes reduced from approximately 180,000 nodes to around 8,000 nodes.

References

- [1] F. Bonnet, J. Mazari, P. Cinnella and P. Gallinari, Airfrans: High fidelity computational fluid dynamics dataset for approximating Reynolds-averaged Navier–Stokes solutions. *Adv. in Neural Inf. Process. Syst.* **35**, 23463–23478 (2022)
- [2] F. Bordeu, F. Casenave and J. Cortial, *BasicTools: a numerical simulation toolbox*. *J. Open Source Softw.* **8**, article no. 5142 (2023)
- [3] F. Casenave and N. Akkari, *An error indicator-based adaptive reduced order model for nonlinear structural mechanics – application to high-pressure turbine blades*. *Math. Comput. Appl.* **24**, article no. 41 (2019)
- [4] F. Casenave, N. Akkari, F. Bordeu, C. Rey and D. Ryckelynck, *A noninvasive distributed reduced-order modeling framework for nonlinear structural mechanics – application to elastoviscoplastic computations*. *Internat. J. Numer. Methods Engrg.* **121**, 32–53 (2020)
- [5] F. Casenave, A. Ern and T. Lelièvre, *Accurate and online-efficient evaluation of the a posteriori error bound in the reduced basis method*. *ESAIM Math. Model. Numer. Anal.* **48**, 207–229 (2014)
- [6] F. Casenave, A. Gariah, C. Rey and F. Feyel, *A noninvasive reduced order model for nonlinear transient thermal problems with non-parametrized variability*. *Adv. Model. Simul. Eng. Sci.* **7**, 1–19 (2020)
- [7] F. Casenave, X. Roynard and B. Staber, *Tensile2d: 2D quasistatic non-linear structural mechanics solutions, under geometrical variations*. URL <https://doi.org/10.5281/zenodo.10124594>, Zenodo (2023)
- [8] F. Casenave, X. Roynard and B. Staber, *Tensile2d: 2D quasistatic non-linear structural mechanics solutions, under geometrical variations*. URL <https://huggingface.co/datasets/PLAID-datasets/Tensile2d>, Hugging Face (2024)
- [9] F. Casenave, B. Staber and X. Roynard, *MMGP: a mesh morphing Gaussian process-based machine learning method for regression of physical problems under non-parametrized geometrical variability*. *Adv. Neural Inf. Process. Syst.* **36** (2024)
- [10] CFD General Notation System (CGNS). <https://cgns.github.io>, accessed: 2024-08-06
- [11] T. Daniel, *Machine learning for nonlinear model order reduction*. Ph.D. thesis, Université Paris sciences et lettres, URL <http://www.theses.fr/2021UPSLM039>, (2021)
- [12] T. Daniel, F. Casenave, N. Akkari, A. Ketata and D. Ryckelynck, *Physics-informed cluster analysis and a priori efficiency criterion for the construction of local reduced-order bases*. *J. Comput. Phys.* **458**, article no. 111120 (2022)
- [13] T. Daniel, F. Casenave, N. Akkari and D. Ryckelynck, *Model order reduction assisted by deep neural networks (ROM-net)*. *Adv. Model. Simul. Eng. Sci.* **7**, article no. 16 (2020)
- [14] T. Daniel, F. Casenave, N. Akkari and D. Ryckelynck, *Data augmentation and feature selection for automatic model recommendation in computational physics*. *Math. Comput. Appl.* **26**, article no. 17 (2021)
- [15] T. Daniel, F. Casenave, N. Akkari, D. Ryckelynck and C. Rey, *Uncertainty quantification for industrial numerical simulation using dictionaries of reduced order models*. *Mech. Ind.* **23**, article no. 3 (2022)
- [16] IRT SystemX et al., *ML4PhySim challenge*. <https://www.codabench.org/competitions/1534>, accessed: 2024-08-07
- [17] IRT SystemX et al., *NeurIPS 2024 – ML4CFD Competition*. <https://www.codabench.org/competitions/3282>, accessed: 2024-08-07
- [18] MOR_DICUS consortium, *Mordicus*. code: <https://gitlab.com/mor-dicus/mordicus>, doc: <https://mordicus.readthedocs.io>, accessed: 2024-08-07
- [19] A. Kabalan, F. Casenave, F. Bordeu, V. Ehrlacher and A. Ern, *Elasticity-based mesh morphing technique with application to reduced-order modeling*. arXiv:2407.02433v2 (2024)
- [20] A. T. Patera and G. Rozza, *Reduced basis approximation and a posteriori error estimation for parametrized partial differential equations*. MIT Pappalardo Graduate Monographs in Mechanical Engineering, MIT, Cambridge, MA (2007)
- [21] X. Roynard, F. Casenave and B. Staber, *Rotor37: a 3D CFD RANS dataset, under geometrical variations of a compressor blade*. URL <https://doi.org/10.5281/zenodo.10149830>, Zenodo (2023)
- [22] X. Roynard, F. Casenave and B. Staber, *Airfrans_original*. Zenodo (2024)
- [23] X. Roynard, F. Casenave and B. Staber, *Airfrans_original*. URL https://huggingface.co/datasets/PLAID-datasets/AirFRANS_original, Hugging Face (2024)
- [24] X. Roynard, F. Casenave and B. Staber, *Rotor37: a 3D CFD RANS dataset, under geometrical variations of a compressor blade*. URL <https://huggingface.co/datasets/PLAID-datasets/Rotor37>, Hugging Face (2024)
- [25] D. Ryckelynck, F. Casenave and N. Akkari, *Resources: Software and tutorials*. In *Manifold Learning: Model Reduction in Engineering*, pp. 53–69, Springer, Cham (2024)
- [26] Safran, *genericROM*. code: <https://gitlab.com/drti/genericrom>, doc: <https://genericrom.readthedocs.io>, accessed: 2024-08-07
- [27] Safran, *Lagun*. code: <https://gitlab.com/drti/lagun>, doc: <https://lagun.readthedocs.io>, accessed: 2024-08-07

- [28] Safran, MMGP. code: <https://gitlab.com/drti/mmgp>, doc: <https://mmgp.readthedocs.io>, accessed: 2024-08-07
- [29] Safran, Muscat. code: <https://gitlab.com/drti/muscat>, doc: <https://muscat.readthedocs.io>, accessed: 2024-08-07
- [30] Safran, PLAID. code: <https://gitlab.com/drti/plaid>, doc: <https://plaid-lib.readthedocs.io>, accessed: 2024-08-07

Fabien Casenave is a research scientist in applied mathematics and a team leader at SafranTech, the research center of the Safran Group, a leading French multinational aerospace and defense corporation. At SafranTech's Digital Sciences and Technologies department, he focuses on developing artificial intelligence and reduced-order modeling technologies to accelerate industrial computations. He was awarded the 2024 Felix Klein Prize.

fabien.casenave@safrangroup.com