CSMA 2026
17ème Colloque National en Calcul des Structures
18-22 Mai 2026, Presqu'île de Giens (Var)

Approach for learning constitutive laws in the context of experimental mechanics

L. Mabileau, C. Jailin, E. Baranger

Université Paris-Saclay, CentraleSupélec, ENS Paris-Saclay, CNRS, LMPS - Laboratoire de Mécanique Paris-Saclay Gif-sur-Yvette
{lucas.mabileau, clement.jailin, emmanuel.baranger}@ens-paris-saclay.fr

Résumé — To design a structure, it is essential to have a reliable constitutive law. When using neural network-based modeling, this reliability depends on the data selected for training, raising the question of which experiments should be performed. Which loadings should be chosen to provide the most informative data ? One approach is to define, a priori, a quality measure based solely on the observations and correlated with the model's performance. This measure can then serve as an indicator to guide the collection of experimental data.
Mots clés — Physics Augmented Neural Networks, Experimental design, Constitutive laws, identification, Machine learning.

1 Introduction

In engineering, the design of a structure relies on a methodical and iterative process comprising multiple stages. Based on the specifications, design constraints, as well as sizing, qualitative or quantitative criteria, are established to ensure the function and performance of the component to be designed. The structural modeling phase occupies a central role in this process, as it enables a rapid and cost-effective assessment of the component's compliance with the selected criteria, while also contributing to the optimization of the overall design process. In cases involving criteria related to mechanical performance, it is essential to adequately represent the geometry, the loading conditions and the behavior of the material used, as the quality of these modeling choices directly affects the relevance of the results obtained.

This is particularly true for the modeling of material constitutive laws. Therefore, for each material, it is necessary, at a minimum, to establish a set of basic experimental tests to calibrate the models employed. The form of these models must be adapted on a case-by-case basis, while remaining consistent with the general framework of the thermodynamics of irreversible processes. Consequently, this modeling activity is a recurrent aspect of an engineer's work.

However, in the most complex cases, determining a suitable mathematical form for a model that accurately captures the actual behavior of the material remains a major challenge. Moreover, over the past decade, machine learning has gained a prominent role in mechanical modelling [1], driven by data-intensive algorithms that can be effectively combined with increasingly rich experimental measurements [2].

The current trend, referred to as data-augmented engineering, aims to integrate well-established physical frameworks directly into neural networks, thereby constraining them to satisfy, for example, thermodynamic principles, while benefiting from greater expressive flexibility than traditional approaches. In the case of modeling a hyperelastic material, a particularly promising architecture is the Physics-Augmented Neural Network (PANN) [3, 4, 5]. The main components of this architecture involve explicitly enforcing material objectivity through an invariant-based description and, more importantly, employing Input Convex Neural Networks (ICNNs) [6] to construct a convex energy potential with respect to the input variables. The resulting stress responses are then obtained via automatic differentiation of the energy.

The use of these new tools raises questions regarding the identification of constitutive laws from experimental data. Indeed, unlike traditional approaches, which are parsimonious in terms of the number of parameters to be identified and whose sensitivities can be controlled and optimized at the scale of a single experiment, neural network-based models contain a large number of trainable parameters, each of which exerts a localized influence on the overall model response. Consequently, confidence in the

model calibration is directly linked to the richness and quality of the data used during the neural network training.

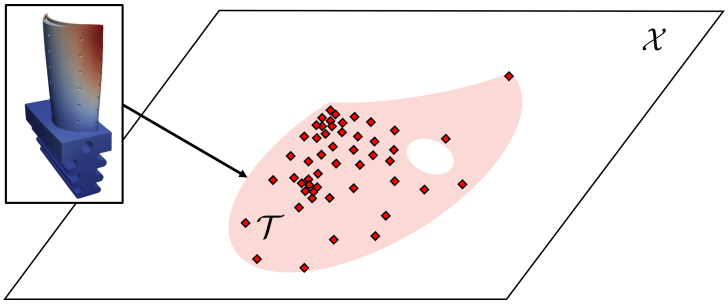

FIGURE 1 – Schematic representation of the kinematic space $\mathcal{X}$, along with the target subspace $\mathcal{T}$ (green area) and the set of points describing this region based on the discretization using a finite element mesh.

Let $\mathcal{X}$ denote the space of mechanical states. This space consists, at a minimum, of observable variables such as kinematic quantities or temperature. In the case of complex behaviors, the dimension of this space increases due to dependence on internal variables, such as plastic strain, or damage variables. If we restrict ourselves to the purely kinematic case, as in hyperelastic behaviors, this space can be described by the Green–Lagrange strain tensor, that is, the space of second-order symmetric tensors, which have six independent components.

In a learning-based identification framework, data drawn from the space $\mathcal{X}$ serve a dual purpose : they are used both to train the model and to evaluate its generalization capability through validation or test sets. The quality and representativeness of these samples directly determine the reliability of the resulting model. This raises the issue of selecting the data necessary for training the corresponding constitutive model. Indeed, it is impossible to explore the entirety of the space $\mathcal{X}$.

In the design process, the validity of the model must be ensured over a usage range corresponding to a specific application, determined by the structure and the expected loadings. The space $\mathcal{X}$ is thus restricted to a relevant subset, referred to as the *target space* $\mathcal{T}$. This space is not chosen arbitrarily : starting from the initial problem and an assumed material behavior, it is possible to estimate the mechanical states that are actually experienced by the structure, as illustrated in Figure 1.

Since the objective is to assess the relevance of a dataset with respect to the design requirements, the target space does not have a uniform importance : certain kinematic states play a more critical role than others in the selected design criteria. Beyond the purely statistical aspect, it may be relevant to highlight specific points in the space that have a particular influence on the design criteria. For instance, one may seek to design the structure based on the maximum stress or a structural quantity, such as the maximum displacement, thereby emphasizing regions of low stiffness and the associated local behavior. The quality of the model training can then be evaluated using *performance metrics* defined over the *target space*.

Alongside the choice of the target space arises the question of its sampling. From an experimental perspective, this involves determining which tests should be performed to meaningfully inform and refine the model. This issue lies at the intersection of the literature on *space-filling* and *design of experiments*, which aims to select the most informative experiments to estimate a model with a limited number of experiments.

To this end, it is necessary to define *information content measures*, serving as *a priori* indicators of the usefulness of the generated data. These measures allow the experimenter to prioritize the most relevant experiments and, ideally, they should be correlated with the previously defined *performance metric* to ensure that the selected experiments effectively contribute to improving the model.

The study that follows constitutes a preliminary investigation of this issue. The initial problem is simplified to facilitate understanding, and a target space is defined, along with an associated performance metric. A data catalog is then generated, from which an *information content measure* is proposed, based solely on the comparison of datasets. Finally, the correlation between this measure and the model's performance over the target space is examined, in order to assess its ability to guide the selection of the most relevant experiments.

2 Problem formulation and data generation

For simplicity in visualizing the phase space, it is reduced to two dimensions in this study. The corresponding representation, shown in Figure 2, is constructed on a regular 64×64 grid. The hidden model examined in this work is defined as follows :

$$\mathcal{X} \to \mathbb{R}, \qquad (x,y) \mapsto z = x^2 + y^2 \tag{1}$$

Although simple, this convex energy serves as a representative analogue of the type of behaviors being learned, while clearly isolating the effects related to the sampling of the phase space. The simplicity of this potential (here quadratic) does not affect the generality of the approach : the network, by virtue of its universal approximation capability, can adapt equally well to a simple energy or to a more complex traditional form.

The main hypothesis adopted here is that the energies z_i, associated with each kinematic measurement point (x_i, y_i), are known *a priori*. Since such energy measurements are not accessible experimentally, this assumption places the study in a supervised learning framework to simplify the analysis. However, numerous unsupervised approaches exist (EUCLID [7], modified Constitutive Relation Error [8], Virtual Field Method [9]) that allow identification of a constitutive model without direct access to these energies, by solving a global mechanical equilibrium. The supervised option chosen here is therefore not restrictive.

The first step consists of artificially generating the target dataset. This dataset is drawn from a Gaussian distribution, consisting of 5000 points, and then truncated in the phase space to limit the sampling region. This target dataset is denoted as

$$\mathcal{T} : \left\{ (x_i^{\mathcal{T}}, y_i^{\mathcal{T}}), z_i^{\mathcal{T}} \right\}_{i=1}^{5000} \tag{2}$$

Next, a catalog of virtual experiments is constructed, consisting of 1000 measurement points each ; that is, a set of datasets generated according to different strategies and probing various regions of the phase space :

$$\mathcal{D}_j : \left\{ (x_i^{\mathcal{D}_j}, y_i^{\mathcal{D}_j}), z_i^{\mathcal{D}_j} \right\}_{i=1}^{1000} \tag{3}$$

The resulting datasets are illustrated in Figure 2. The objective is to classify these datasets and define criteria to quantify their relative relevance with respect to the target dataset. In this framework, the experimental datasets serve to train the model, while the target data constitute a test set used to evaluate the model's ability to generalize to the intended application. The target dataset thus represents the usage domain over which the model's performance must be guaranteed.

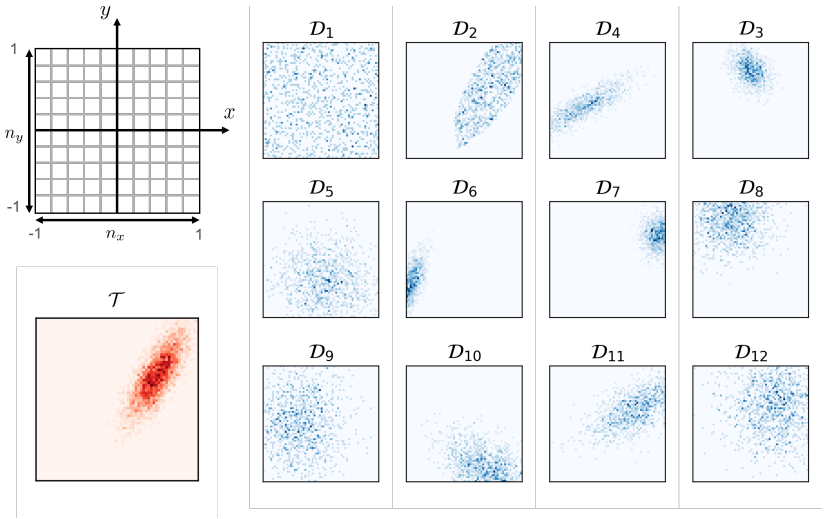

FIGURE 2 – Top left : representation of the phase space $\mathcal{X}$ for the studied problem, discretized on a uniform grid of size $n_x \times n_y = 64 \times 64$, bounded between -1 and 1 in each direction. Bottom left, in red : distribution of the target dataset $\mathcal{T}$, drawn from a truncated Gaussian distribution (*5000 points*) in $\mathcal{X}$. Right, in blue : a set of 12 datasets of *1000 points* each.

3 Method

The first step of the study consists of training, from each dataset $\mathcal{D}_j$, an instance of the model $h_{\mathcal{D}_j}$ as illustrated in Figure 3.

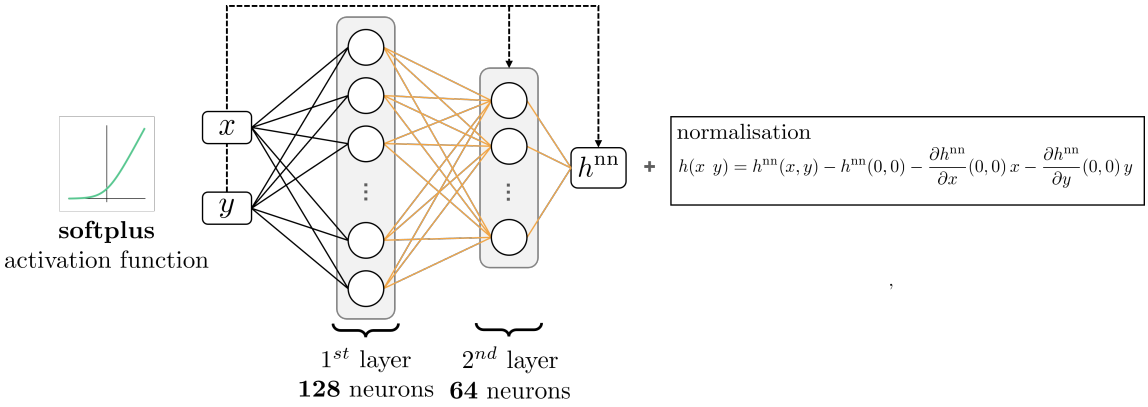

FIGURE 3 – Architecture of the ICNN used, composed of two hidden layers (128 and 64 neurons). The convexity of the potential is enforced through the use of *softplus* activation functions and positive weights (orange lines) in the second and third layers. The use of *skip connections* (dashed arrows) allows information from the input data to propagate through the different layers of the network.

Model This model adopts the architecture of a PANN [4] : it is an *Input Convex Neural Network* (ICNN), in which the convexity of the potential z with respect to the variables (x,y) is enforced by construction. This property is achieved through the use of positive weights in the deeper layers and *softplus* activation functions (convex and non-decreasing). The network consists of two hidden layers of size 128 and 64, with a total of 8,837 trainable parameters.

Furthermore, a normalization step enforces, by construction, that the potential and its gradient (the stresses) vanish when the strain is zero : $z(x=\varepsilon_{11}=0, y=\varepsilon_{22}=0)=0$, $\frac{\partial z}{\partial x}(0,0)=0$, and $\frac{\partial z}{\partial y}(0,0)=0$. Although deliberately simple, this model nevertheless respects all the standard principles of PANNs developed in the literature and experimentally applied to learning real hyperelastic behaviors [10].

This model is trained in a supervised manner on the catalog datasets ($\mathcal{D}_1$ to $\mathcal{D}_{12}$). An additional model is also trained on the target dataset : this does not constitute a practically usable model, since these test data would not be experimentally accessible, but serves solely as a reference model (not necessarily optimal) for evaluating the other models.

Training metric The loss function to be minimized for training the model weights is the mean squared error (MSE), defined as following.

$$\mathcal{L}(\mathcal{D}_j) = \mathrm{MSE}\left(\{z_i^{\mathcal{D}_j}\}_{i=1}^{n_{\mathcal{D}_j}}, \{h_{\mathcal{D}_j}(x_i^{\mathcal{D}_j}, y_i^{\mathcal{D}_j})\}_{i=1}^{n_{\mathcal{D}_j}}\right) = \frac{1}{n_{\mathcal{D}_j}} \sum_{i=1}^{n_{\mathcal{D}_j}} \left(z_i^{\mathcal{D}_j} - h_{\mathcal{D}_j}(x_i^{\mathcal{D}_j}, y_i^{\mathcal{D}_j})\right)^2 \tag{4}$$

where $n_{\mathcal{D}_j}$ denotes the number of points in the dataset $\mathcal{D}_j$, and $h_{\mathcal{D}_j}(x_i^{\mathcal{D}_j}, y_i^{\mathcal{D}_j})$ is the prediction of the model trained on $\mathcal{D}_j$ for the point $(x_i^{\mathcal{D}_j}, y_i^{\mathcal{D}_j})$.

Performance metric The quantity of interest here is the performance of the model calibrated on the target data, which is chosen to be quantified through the evaluation of the MSE :

$$\Phi_{\mathcal{D}_j}(\mathcal{T}) = \mathrm{MSE}\left(\{z_i^{\mathcal{T}}\}_{i=1}^{n_{\mathcal{T}}}, \{h_{\mathcal{D}_j}(x_i^{\mathcal{T}}, y_i^{\mathcal{T}})\}_{i=1}^{n_{\mathcal{T}}}\right) \tag{5}$$

where $n_{\mathcal{T}}$ denotes the number of points in the target dataset $\mathcal{T}$, and $h_{\mathcal{D}_j}(x_i^{\mathcal{T}}, y_i^{\mathcal{T}})$ is the prediction of the model trained on $\mathcal{D}_j$ for the point $(x_i^{\mathcal{T}}, y_i^{\mathcal{T}})$ belonging to the target dataset. This performance metric of the datasets constitutes an *a posteriori* evaluation (performed after model training), and its computation requires knowing the ground truth.

Information content measure In this context, *a priori* measures are introduced, based on the comparison of statistics between two point clouds in the initial space $\mathcal{X}$, intended to be correlated with the dataset performance metric. One measure proposed here corresponds to the inclusion rate of the target data within the training data, that is, an estimate of the average degree to which the former is covered by the latter. To do this, the convex hull of each 2D point cloud is computed. The measure is then defined as the ratio between the "area" of the intersection of the convex hulls (in purple) and the "area" of the convex hull of cloud $\mathcal{A}$ (in blue).

The method for computing this inclusion rate is illustrated in Figure 4.

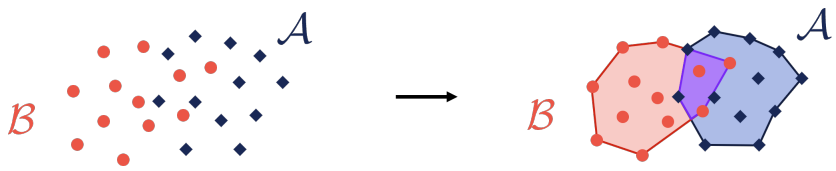

FIGURE 4 – Example of a measure between two point clouds : here, the goal is to evaluate the inclusion rate of point cloud $\mathcal{A}$ within point cloud $\mathcal{B}$.

The convex hull is extracted for each of the two point clouds, allowing the definition of "areas." The inclusion rate is then defined as :

$$\Psi(\mathcal{D}_j, \mathcal{T}) = \frac{S_{\mathcal{D}_j \cap \mathcal{T}}}{S_{\mathcal{T}}} \in [0,1] \tag{6}$$

A value of $\Psi = 0$ indicates that no portion of the target space is covered by the convex hull of the training data, whereas a rate of $\Psi = 1$ signifies that the entire target space is fully contained within this hull.

4 Results and Discussion

Model training The optimizer used is *ADAM*, with a learning rate set to 5×10^{-4} and a batch size of 100, and training is carried out until convergence. An 80/20 split between training and validation data was adopted for each dataset. These optimization parameters were not subject to extensive tuning, as robust convergence was achieved without *fine-tuning*, owing to the model architecture and the low dimensionality of the problem.

The biases of the first layer were initialized to zero, while the weights and remaining biases were initialized according to a Kaiming uniform distribution and subsequently truncated to positive values on the relevant layers to satisfy the ICNN constraints. The absolute performance of the trained models is presented in Table 1.

TABLE 1 – MSE values for each trained model, computed on the training data (first row), on the target data (second row), and the corresponding ratio (third row).

Model	$h_{\mathcal{D}_1}$	$h_{\mathcal{D}_2}$	$h_{\mathcal{D}_3}$	$h_{\mathcal{D}_4}$	$h_{\mathcal{D}_5}$	$h_{\mathcal{D}_6}$
$\mathcal{L}(\mathcal{D}_j)$	8.49×10^{-7}	5.94×10^{-6}	1.27×10^{-6}	2.64×10^{-7}	5.45×10^{-7}	2.85×10^{-7}
$\Phi_{\mathcal{D}_j}(\mathcal{T})$	1.44×10^{-6}	8.23×10^{-6}	1.73×10^{-3}	6.50×10^{-4}	1.14×10^{-4}	3.95×10^{-2}
Ratio	1.70	1.38	1.36×10^{3}	2.47×10^{3}	2.09×10^{2}	1.38×10^{5}

Model	$h_{\mathcal{D}_7}$	$h_{\mathcal{D}_8}$	$h_{\mathcal{D}_9}$	$h_{\mathcal{D}_{10}}$	$h_{\mathcal{D}_{11}}$	$h_{\mathcal{D}_{12}}$
$\mathcal{L}(\mathcal{D}_j)$	2.85×10^{-7}	2.80×10^{-6}	8.61×10^{-7}	4.06×10^{-6}	3.25×10^{-6}	2.73×10^{-6}
$\Phi_{\mathcal{D}_j}(\mathcal{T})$	1.75×10^{-4}	1.59×10^{-4}	3.97×10^{-5}	3.45×10^{-2}	5.68×10^{-6}	1.65×10^{-6}
Ratio	6.15×10^{2}	5.67×10^{1}	4.61×10^{1}	8.49×10^{3}	1.75	6.05×10^{-1}

By examining the MSE ratios, four datasets exhibit a ratio close to or below 1 : $\mathcal{D}_1$, $\mathcal{D}_2$, $\mathcal{D}_3$, and $\mathcal{D}_4$. This value is interpreted as a measure of the model's ability to perform on the target data relative to what it has learned during training. For example, the raw performance on the target data is better for the model trained on $\mathcal{D}_1$ than for that on $\mathcal{D}_{12}$, although the ratio favors $\mathcal{D}_{12}$, since the model tends to overperform relative to the training data. Indeed, the distributions of the two datasets differ : those of $\mathcal{D}_{12}$ are more concentrated, whereas those of $\mathcal{D}_1$ are more spread out. In the case of $\mathcal{D}_{12}$, the model represents the tails of the distribution less accurately, but it performs well at the density peak, where the target data are also concentrated.

On the computation of the information content measure The inclusion rates are computed based on an estimate of the areas of the intersection and of the convex hull of the target data, obtained using a statistical method that counts the number of "cells" falling within the hull over 10,000 samples. The value of this rate, expressed as a percentage, is presented in Table 2.

TABLE 2 – Inclusion rate of the target data for each dataset in the catalog.

Dataset	$\mathcal{D}_1$	$\mathcal{D}_2$	$\mathcal{D}_3$	$\mathcal{D}_4$	$\mathcal{D}_5$	$\mathcal{D}_6$	$\mathcal{D}_7$	$\mathcal{D}_8$	$\mathcal{D}_9$	$\mathcal{D}_{10}$	$\mathcal{D}_{11}$	$\mathcal{D}_{12}$
Ψ (%)	99.1	97.4	22.4	31.5	77.4	0.0	37.8	61.4	81.0	41.0	95.0	99.8

Among the datasets in the catalog, four exhibit a very high inclusion rate (greater than 95%) : $\mathcal{D}_1$, $\mathcal{D}_2$, $\mathcal{D}_{11}$, and $\mathcal{D}_{12}$. In contrast, dataset $\mathcal{D}_6$ includes no points (0%). It appears that the datasets with a high inclusion rate are also those that demonstrate better performance on the target data.

On the correlation with the performance metric To visually compare the *performance metric* with the inclusion rate measure, these two indicators are plotted together in Figure 5.

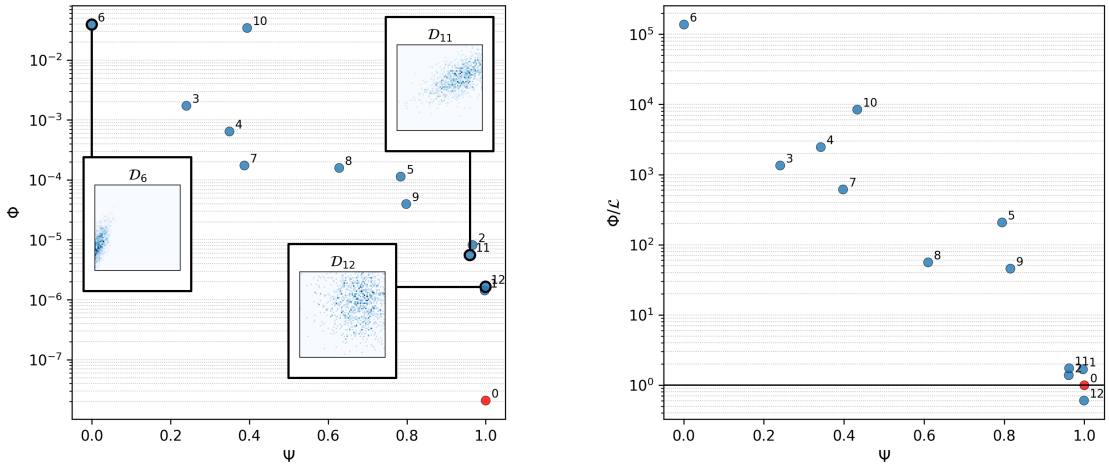

FIGURE 5 – plot : inclusion rate Ψ on the x-axis and performance metric Φ on the y-axis. Right plot : same x-axis Ψ, but the performance metric Φ is normalized by the value of the loss function $\mathcal{L}$ during model training. Both plots use a logarithmic scale on the y-axis. Each point represents a dataset from the catalog.

Consistent with previous observations, the points corresponding to datasets $\mathcal{D}_1$, $\mathcal{D}_2$, $\mathcal{D}_{11}$, and $\mathcal{D}_{12}$ are clustered near the reference point (in red), which corresponds to the evaluation of a model trained directly on the target data. In contrast, dataset $\mathcal{D}_6$ exhibits both the lowest inclusion rate and the poorest performance on the target data. Between these two extremes, a mixed set of the other datasets is observed, within which a monotonic trend linking the inclusion rate to model performance emerges. This correlation can be explained by the simplicity of the latent behavior and the convexity constraints imposed by the ICNN. Since the underlying potential is smooth and strictly convex, the model fails by

extrapolation outside the sampled region, so datasets whose convex hull closely covers the target space naturally exhibit better generalization.

It appears that, for this choice of target data, the inclusion rate constitutes a reliable *information content measure*, as it correlates with the model's performance on the target data. Thus, it would be possible to predict, prior to calibration, the relevance of a dataset or experimental design.

Furthermore, this measure can be used proactively to guide the data acquisition strategy. Indeed, it is possible to generate or select measurement points in such a way as to achieve an inclusion rate as close as possible to 100%, thereby ensuring that the target space is fully covered by the collected experimental data. Such an approach provides an operational criterion for directing the design of experiments toward the most informative regions.

5 Conclusion

The objective of this study was to propose an experimental sampling strategy for the calibration of a highly parameterized neural network-based constitutive model. After formalizing the problem through the definition of a target space to be explored, a simplified case was analyzed to evaluate a measure capable of identifying, within a catalog, the dataset most suitable for training a regularized model representative of the state of the art in constitutive law modeling via neural networks. The inclusion rate was thus introduced as an *information content measure* and demonstrated a strong ability to correlate with the *performance metric* defined on the target data, making it a relevant candidate to guide an effective sampling strategy.

Several natural extensions were then identified, and some preliminary results already suggest promising perspectives. One direction aims to enrich the *information content measure* between datasets by incorporating, beyond the geometry of their convex hull, their local density, in order to better capture the effective distribution of the accessed states. Statistical analysis could also be strengthened by increasing the number of case studies and random draws, which would help consolidate the robustness of the observed trends. Another perspective consists of defining an *information content measure* at the level of each experimental point, allowing the identification and prioritization of the most informative data for the training process. Finally, the diversity of possible measures to compare two datasets suggests an additional avenue for investigation : learning, via neural networks, an *information content measure* optimized to maximize its correlation with the *performance metric*. The latter aspect will be further explored and detailed during the conference.

These various approaches will then need to be challenged in configurations more representative of the experimental context, notably by learning directly from stress data or in an unsupervised framework, thereby opening the way to broader validations and a better understanding of the potential of these sampling strategies in realistic scenarios.

Références

[1] J.N. Fugh, G. Anantha Padmanabha, N. Bouklas, B. Bahmani, W. Sun, N.N. Vlassis, M. Flaschel, P. Carrara, L. De Lorenzis. *A Review on Data-Driven Constitutive Laws for Solids*, Archives of Computational Methods in Engineering, Springer, 1-43, 2024.

[2] F. Hild, S. Roux. *Digital Image Correlation : from Displacement Measurement to Identification of Elastic Properties – a Review*, Strain, Wiley Online Library, 69-80, 2006.

[3] F. As'ad, P. Avery, C. Farhat. *A mechanics-informed artificial neural network approach in data-driven constitutive modeling*, International Journal for Numerical Methods in Engineering, Wiley Online Library, 2738-2759, 2022.

[4] L. Linden, D.K. Klein, K.A. Kalina, J. Brummund, O. Weeger, M. Kästner. *Neural networks meet hyperelasticity : A guide to enforcing physics*, Journal of the Mechanics and Physics of Solids, ScienceDirect, 129-105363, 2023.

[5] K.A. Kalina, J. Brummund, W. Sun, M. Kästner. *Neural networks meet anisotropic hyperelasticity : A framework based on generalized structure tensors and isotropic tensor functions*, Computer Methods in Applied Mechanics and Engineering, Wiley Online Library, 117725, 2025.

[6] B. Amos, L. Xu, J.Z. Kolter. *Input Convex Neural Networks*, PMLR, 146-155, 2017.

[7] P. Thakolkaran, A. Joshi, Y. Zheng, M. Flaschel, L. De Lorenzis, S. Kumar. *NN-EUCLID : Deep-learning hyperelasticity without stress data*, Journal of the Mechanics and Physics of Solids, ScienceDirect, 105076, 2022.

[8] A. Benady, E. Baranger, L. Chamoin. *NN-mCRE : A modified constitutive relation error framework for unsupervised learning of nonlinear state laws with physics-augmented neural networks*, International Journal for Numerical Methods in Engineering, Wiley Online Library, e7439, 2024.

[9] R. Lourenço, P. Georgieva, E. Cueto, A. Andrade-Campos. *An indirect training approach for implicit constitutive modelling using recurrent neural networks and the virtual fields method*, Computer Methods in Applied Mechanics and Engineering, ScienceDirect, 116961, 2024.

[10] C. Jailin, A. Benady, R. Legroux, E. Baranger. *Experimental Learning of a Hyperelastic Behavior with a Physics-Augmented Neural Network*, Experimental Mechanics, Springer, 1465-1481, 2024.