Enforcing physics onto PINNs for more accurate inhomogeneous material identification

van der Heijden, B.; Li, X.; Lubineau, G.; Florentin, E.

DOI
10.1016/j.cma.2025.117993

Publication date
2025

Document Version
Final published version

Published in
Computer Methods in Applied Mechanics and Engineering

Citation (APA)
van der Heijden, B., Li, X., Lubineau, G., & Florentin, E. (2025). Enforcing physics onto PINNs for more accurate inhomogeneous material identification. *Computer Methods in Applied Mechanics and Engineering*, *441*, Article 117993. https://doi.org/10.1016/j.cma.2025.117993

Computer Methods in Applied Mechanics and Engineering 441 (2025) 117993

Contents lists available at ScienceDirect

Comput. Methods Appl. Mech. Engrg.

journal homepage: www.elsevier.com/locate/cma

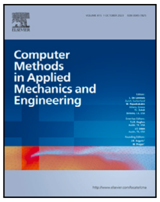

Enforcing physics onto PINNs for more accurate inhomogeneous material identification

B. van der Heijden [a,b,*], X. Li [a], G. Lubineau [a], E. Florentin [c]

[a] King Abdullah University of Science and Technology, Mechanical Engineering Program, Physical Sciences and Engineering Division, Mechanics of Composites for Energy and Mobility Lab., Thuwal, 23955-6900, Saudi Arabia
[b] Delft University of Technology, Maritime and Transport Technology Department, Delft, 2628 CD, The Netherlands
[c] INSA Centre Val de Loire, Univ. Orléans, Univ. Tours, Laboratoire de Mécanique Gabriel LaMé, Bourges, 18022 CEDEX, France

ARTICLE INFO

Keywords:
Physics-informed neural networks
Inverse methods
Material identification
Constitutive equation gap method
Differentiable programming

ABSTRACT

Physics-Informed Neural Networks (PINNs) are computationally efficient tools for addressing inverse problems in solid mechanics, but often face accuracy limitations when compared to traditional methods. We introduce a refined PINN approach that rigorously enforces certain physics constraints, improving accuracy while retaining the computational benefits of PINNs. Unlike conventional PINNs, which are trained to approximate (differential) equations, this method incorporates classical techniques, such as stress potentials, to satisfy certain physical laws. The result is a physics-enforced PINN that combines the precision of the Constitutive Equation Gap Method (CEGM) with the automatic differentiation and optimization frameworks characteristic of PINNs. Numerical comparisons reveal that the enforced PINN approach indeed achieves near-CEGM accuracy while preserving the efficiency advantages of PINNs. Validation through real experimental data demonstrates the ability of the method to accurately identify material properties and inclusion geometries in inhomogeneous samples.

1. Introduction

In recent years, significant progress in solid mechanics has been driven by the pursuit of precise predictive methods in engineering and science. These advances highlight the pivotal role of modelling in understanding material behaviour, optimizing mechanical systems, and designing structures. While the fundamentals of solid mechanics provide a basis for kinematic and static admissibility, challenges persist in formulating constitutive equations [1]. Despite recent developments [2], a persistent challenge lies in the precise determination and application of constitutive equations, particularly in complex systems with heterogeneities. The process of conducting material characterization experiments, extracting properties, and translating findings into accurate constitutive equations demands considerable expertise. This paper addresses the determination of constitutive properties in inhomogeneous structures by combining traditional inverse methods with machine learning technologies.

The traditional approach to material identification relies on simple tests, assuming homogeneity in stress and strain fields for analytical analysis and property extraction. However, this assumption is violated upon approaching failure, as non-homogeneous stress–strain components emerge, even for simple examples such as necking-localization during tensile tests. Moreover, this approach is time-consuming and expensive, as only few properties can be obtained simultaneously, and thus multiple tests are required to characterize a single material completely.

---

[*] Corresponding author at: Delft University of Technology, Maritime and Transport Technology Department, Delft, 2628 CD, The Netherlands.
E-mail addresses: B.vanderHeijden@tudelft.nl (B. van der Heijden), Eric.Florentin@insa-cvl.fr (E. Florentin).

https://doi.org/10.1016/j.cma.2025.117993
Received 26 November 2024; Received in revised form 1 April 2025; Accepted 2 April 2025

Modern approaches employ full-field measurement techniques, such as digital image correlation (DIC), under more complex loading conditions. While this allows for the simultaneous extraction of many material properties, solving the resulting inverse problem requires more sophisticated and numerical methods. Examples include the Finite Element Model Updating (FEMU) [3,4], the Equilibrium Gap Method (EGM) [5], the Virtual Fields Method (VFM) [6,7], and Constitutive Equation Gap Method (CEGM) [8]. For a comprehensive understanding and comparison of these methodologies, readers can refer to existing literature [9,10]. These methods become computationally expensive, especially when dealing with inhomogeneous material properties [5,11].

Recent advances, such as Physics Informed Neural Networks (PINNs), offer promising alternatives [12]. PINNs, partially trained on information from differential equations, prove effective for both forward and inverse problems [13], enabling efficient identification of material parameters in inhomogeneous solids [14]. Wei et al. [15] explored a novel approach that combines the CEGM formulation with PINNs, resulting in an inverse method which is orders of magnitude faster than the original CEGM method. The increased speed is attributed to Automatic Differentiation (AD) and the parallelization inherently available in machine learning frameworks. However, the PINN approaches generally sacrifice accuracy in recovered material properties [15].

This paper addresses the accuracy limitation by reintroducing traditional physics and numerical strategies into the design of NNs and PINNs. It will show that designing PINNs that enforce physics, rather than informing the network, will improve performance. Section 2 introduces the inverse problem and existing solving strategies, outlining the notation and definitions used further on. Subsequently, Section 3 introduces various strategies of enforcing certain physics rules, rather than teaching them. These include thermodynamic admissibility of the constitutive equations and impose equilibrated stress fields. Numerical examples are discussed in Section 4 to illustrate the benefits and drawbacks of these techniques. Then, Section 5 highlights the applicability of the inverse strategy with an experimental example. Lastly, the research is concluded in Section 6.

2. Inverse problems and the state of the art

In the context of an inverse problem, the objective is to ascertain the material properties of a given structure by utilizing measured data. In contrast to conventional forward problems, where known inputs are employed to predict behaviour, inverse problems commence with observed behaviour and attempt to reverse-engineer the responsible material properties.

Nevertheless, the model relies on the same fundamentals. Here a local-continuum formulation with Cauchy stress is assumed within the regime of small perturbations. The solution to a forward problem adheres to the following criteria:

- Kinematic admissibility:

$$\begin{cases} \underline{\underline{\varepsilon}} = \frac{1}{2}\left(\underline{u} \otimes \underline{\boldsymbol{\nabla}} + \underline{\boldsymbol{\nabla}} \otimes \underline{u}\right) & \forall \underline{x} \in \Omega \\ \underline{u} = \underline{\bar{u}} & \forall \underline{x} \in \partial\Omega_{\underline{u}} \end{cases} \tag{1}$$

- Static admissibility:

$$\begin{cases} \underline{0} = \underline{\boldsymbol{\nabla}} \cdot \underline{\underline{\sigma}} & \forall \underline{x} \in \Omega \\ \underline{\underline{\sigma}} = \underline{\underline{\sigma}}^T & \forall \underline{x} \in \Omega \\ \underline{\bar{t}} = \underline{\underline{\sigma}} \cdot \underline{n} & \forall \underline{x} \in \partial\Omega_{\underline{t}} \end{cases} \tag{2}$$

- Constitutive equation for this publication is assumed to be linear elastic:

$$\left\{ \underline{\underline{\sigma}} = \underline{\underline{\underline{\underline{A}}}} : \underline{\underline{\varepsilon}} \qquad \forall \underline{x} \in \Omega \right. \tag{3}$$

Here, Ω represents the problem domain, $\underline{u}$ displacement, $\underline{\underline{\varepsilon}}$ infinitesimal strain, $\underline{\underline{\sigma}}$ Cauchy stress, and $\underline{\underline{\underline{\underline{A}}}}$ the stiffness tensor, which can vary spatially. The domain boundary, $\partial\Omega$, is divided into kinematic and static regions, denoted as $\partial\Omega_{\underline{u}}$ and $\partial\Omega_{\underline{t}}$, respectively, with no overlap between them. Along the kinematic boundary, displacement $\underline{\bar{u}}$ is specified, while the static boundary has prescribed traction $\underline{\bar{t}}$, as illustrated in Fig. 1.

To facilitate the description of these fields, the following tensor spaces are introduced. All kinematically admissible displacement fields are contained within:

$$\mathcal{K}(\underline{\bar{u}}) = \left\{ \underline{u} \in \mathcal{V}_u \mid \; \underline{u} = \underline{\bar{u}}\,, \; \forall \underline{x} \in \partial\Omega_{\underline{u}} \right\}, \tag{4}$$

where $\mathcal{V}_u$ represents the vector field that is in H^1. Similarly, all statically admissible stress fields are encompassed by:

$$\mathcal{S}(\underline{\bar{t}}) = \left\{ \underline{\underline{\sigma}} \in \mathcal{V}_\sigma \mid \; \underline{\boldsymbol{\nabla}} \cdot \underline{\underline{\sigma}} = \underline{0}\,, \; \forall \underline{x} \in \Omega \text{ and } \underline{\underline{\sigma}} \cdot \underline{n} = \underline{\bar{t}}\,, \; \forall \underline{x} \in \partial\Omega_{\underline{t}} \right\}, \tag{5}$$

where $\mathcal{V}_\sigma$ is the space of all symmetric square-integrable second-order tensor fields. Finally, a space of thermodynamically admissible stiffness tensors exists:

$$\mathcal{A} = \left\{ \underline{\underline{\underline{\underline{A}}}} \in \mathcal{V}_A \mid \; \underline{\underline{\varepsilon}} : \underline{\underline{\underline{\underline{A}}}} : \underline{\underline{\varepsilon}} \geq 0\,, \; \forall \underline{\underline{\varepsilon}} \in \mathcal{V}_\sigma \text{ and } \underline{\underline{\varepsilon}} = \underline{\underline{0}} \iff \underline{\underline{\varepsilon}} : \underline{\underline{\underline{\underline{A}}}} : \underline{\underline{\varepsilon}} = 0 \right\}, \tag{6}$$

where $\mathcal{V}_A$ represents the space of positive definite fourth-order tensor fields with the appropriate symmetries. In the case of heterogeneous, isotropic, and elastic materials, $\mathcal{V}_A$ can be described by two space-dependent parameters, such as Young's modulus $E(\underline{x})$ and Poisson's ratio $\nu(\underline{x})$.

In the inverse problem, the displacement field $\underline{\hat{u}}$ is known, as it was measured with DIC/DVC. The solution concerns finding the constitutive field $\underline{\underline{\underline{\underline{A}}}} \in \mathcal{A}$ and a stress field $\underline{\underline{\sigma}} \in \mathcal{S}(\underline{\bar{t}})$, such that the constitutive requirements, see Eq. (3), are satisfied.

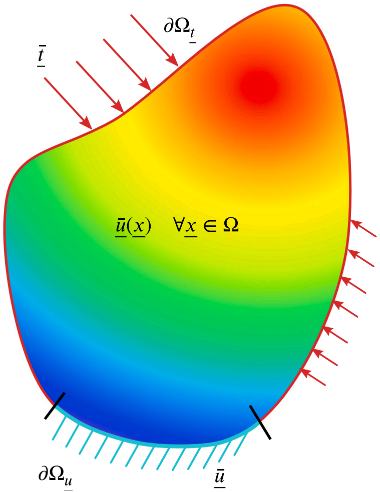

Fig. 1. Example inverse problem consisting of the measured displacement field $\underline{\hat{u}}$, the boundary displacement $\underline{\bar{u}}$ and traction $\underline{\bar{t}}$.

2.1. Constitutive equation gap method

An elegant formulation to quantify the constitutive error relates the measured displacement field to a proposed and a statically admissible stress field through the constitutive equation gap:

$$\mathcal{L}(\underline{\underline{\sigma}}, \underline{\underline{\underline{\underline{A}}}}) = \frac{1}{2} \int_{\Omega} (\underline{\underline{\sigma}} - \underline{\underline{\underline{\underline{A}}}} : \underline{\underline{\varepsilon}}[\underline{\hat{u}}]) : \underline{\underline{\underline{\underline{A}}}}^{-1} : (\underline{\underline{\sigma}} - \underline{\underline{\underline{\underline{A}}}} : \underline{\underline{\varepsilon}}[\underline{\hat{u}}])\, \mathrm{d}\underline{x} \quad \forall \underline{\underline{\sigma}} \in \mathcal{S},\ \underline{\underline{\underline{\underline{A}}}} \in \mathcal{A}, \tag{7}$$

which leverages the positive definiteness of the stiffness tensor to establish an energy norm. Minimizing this energy norm allows for the recovery of the stress and constitutive fields [11,16–18].

Inaccurate predictions might stem from noisy measurements or inadmissible displacement fields. In those cases, a modified formulation is created, introducing an alternative displacement field $\underline{u}$ [19]. This displacement field should be kinematically admissible, but also as close as possible to the measured field. This results in the following penalized minimization problem,

$$\underline{u}^{\star}, \underline{\underline{\sigma}}^{\star}, \underline{\underline{\underline{\underline{A}}}}^{\star} = \underset{\underline{u}\in\mathcal{K},\ \underline{\underline{\sigma}}\in\mathcal{S},\ \underline{\underline{\underline{\underline{A}}}}\in\mathcal{A}}{\operatorname{argmin}} \mathcal{L}(\underline{u}, \underline{\underline{\sigma}}, \underline{\underline{\underline{\underline{A}}}}) + \frac{\alpha}{2} \int_{\Omega} \left\| \underline{u} - \underline{\hat{u}} \right\|^2 \mathrm{d}\underline{x}, \tag{8}$$

where α is a hyperparameter that should be selected. Details on the selection of α can be found in [20,21] and automated selection of the parameter is proposed by Benady et al. [22].

Some of these modified formulations even assume inaccuracies regarding static boundary conditions [20]. They identify that the traction applied or measured along a boundary might include inaccuracies, and add them as penalization terms as well. However, all methods assume that internal static admissibility needs to be satisfied exactly, and most also assume that traction-free boundaries are reliable. As a result, these stay enforced or constrained, rather than penalized.

A drawback of inverse CEGM methods is the complexity of the optimization. This can be resolved by searching through the static and constitutive space separately in a nested optimization. For example, an outer loop attempts to determine the stiffness tensor while the inner loop finds the optimal stress field that belongs to that stiffness. This approach makes the optimization expensive, as every major iteration needs to solve a normal forward problem. Progress has been made, and more efficient strategies exist [16,18]. Nevertheless, finding the constitutive behaviour of inhomogeneous materials remains expensive [10,15].

2.2. Inverse methods with PINNs

PINNs offer a computationally efficient approach for solving inverse problems in solid mechanics, building upon the foundational work by Raissi et al. [13]. PINNs are a class of neural networks that leverage both observed data and the governing differential equations to find solutions to complex physical problems. Wei et al. [15] developed a PINN to identify material properties in inhomogeneous elastic materials. What follows here is a summary of this approach.

The flowchart of Fig. 2 shows how the PINN employs NNs to represent the displacement field, the Cauchy stress field, and the constitutive parameters field. These NNs, combined with automatic differentiation, enable the calculation of internal equilibrium ($\underline{\nabla} \cdot \underline{\underline{\tilde{\sigma}}}$) and the deformation gradient ($\underline{\tilde{u}} \otimes \underline{\nabla}$). The core idea is to minimize the sum of various loss terms, representing the different physics requirements mentioned above, to solve the inverse problem efficiently.

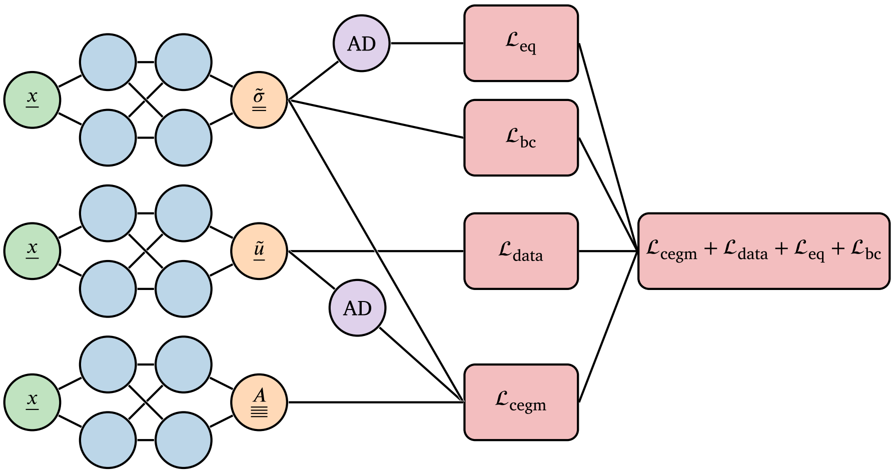

Fig. 2. Three neural networks represent, from top to bottom, the static, kinematic, and constitutive fields of the problem. Automatic differentiation (AD) is used to compute internal equilibrium ($\underline{\nabla} \cdot \underline{\underline{\tilde{\sigma}}}$) and the deformation gradient ($\underline{\tilde{u}} \otimes \underline{\nabla}$). Minimizing the sum of losses solves the inverse problem.

This PINN-based inverse method augments the limited experimental data with the governing differential equations. The loss terms used in this approach include:

- **Data Loss ($\mathcal{L}_{\text{data}}$):** quantifies the discrepancy between the NN displacement $\underline{\tilde{u}}$ and the measured displacement data $\underline{\hat{u}}$ at locations $\underline{x}_j$ where measurements were conducted. ΔV_i is the unit volume to ensure that this represents an integration.

$$\mathcal{L}_{\text{data}} = \sum_{i=1}^{N_u} \left\| \underline{\tilde{u}}(\underline{x}_i) - \underline{\hat{u}}(\underline{x}_i) \right\|^2 \Delta V_i. \tag{9}$$

- **Equilibrium Loss ($\mathcal{L}_{\text{eq}}$):** evaluates whether NN stress field $\underline{\underline{\tilde{\sigma}}}$ satisfies the internal equilibrium equation. This function can be evaluated at any point. Here, a dense sampling of $\underline{x}_j$ is used.

$$\mathcal{L}_{\text{eq}} = \sum_{j=1}^{N_e} \left\| \underline{\nabla} \cdot \underline{\underline{\tilde{\sigma}}}(\underline{x}_j) \right\|^2 \Delta V_j. \tag{10}$$

- **Boundary Condition Loss ($\mathcal{L}_{\text{bc}}$):** evaluates the violation of the static boundary conditions by comparing the NN stresses along the boundary coordinates $\underline{x}_k$ to the traction prescribed on these boundaries $\underline{\bar{t}}$. Here, ΔS_k is the unit surface of the integration.

$$\mathcal{L}_{\text{bc}} = \sum_{k=1}^{N_b} \left\| \underline{\underline{\tilde{\sigma}}}(\underline{x}_k) \cdot \underline{n}(\underline{x}_k) - \underline{\bar{t}}(\underline{x}_k) \right\|^2 \Delta S_k. \tag{11}$$

- **Constitutive Equation Loss ($\mathcal{L}_{\text{cegm}}$):** evaluates the constitutive equation gap, as introduced in the CEGM approach above. The function is evaluated at the same locations $\underline{x}_j$ as the equilibrium loss.

$$\mathcal{L}_{\text{cegm}} = \frac{1}{2} \sum_{j=0}^{N_a} \left(\underline{\underline{\tilde{\sigma}}}(\underline{x}_j) - \underline{\underline{\underline{\underline{\tilde{A}}}}}(\underline{x}_j) : \underline{\underline{\varepsilon}}\left[\underline{\tilde{u}}(\underline{x}_j)\right] \right) : \underline{\underline{\underline{\underline{\tilde{A}}}}}^{-1}(\underline{x}_j) : \left(\underline{\underline{\tilde{\sigma}}}(\underline{x}_j) - \underline{\underline{\underline{\underline{\tilde{A}}}}}(\underline{x}_j) : \underline{\underline{\varepsilon}}\left[\underline{\tilde{u}}(\underline{x}_j)\right] \right) \Delta V_j. \tag{12}$$

Eventually the final solution is found in the following minimization problem:

$$\underline{\tilde{u}}^{\star}, \underline{\underline{\tilde{\sigma}}}^{\star}, \underline{\underline{\underline{\underline{\tilde{A}}}}}^{\star} = \underset{\underline{\tilde{u}} \in \mathcal{N}^u, \underline{\underline{\tilde{\sigma}}} \in \mathcal{N}^\sigma, \underline{\underline{\underline{\underline{\tilde{A}}}}} \in \mathcal{N}^A}{\operatorname{argmin}} \mathcal{L}_{\text{cegm}} + \alpha \mathcal{L}_{\text{data}} + \beta \mathcal{L}_{\text{eq}} + \gamma \mathcal{L}_{\text{bc}}, \tag{13}$$

where $\mathcal{N}^u$, $\mathcal{N}^\sigma$, and $\mathcal{N}^A$ are the spaces spanned by the neural networks, while α, β, and γ represent the penalization factors.

The PINN-based inverse method has the advantage of computational efficiency, reducing the time from hours to just minutes. This improved efficiency results from two key factors. Firstly, the absence of nested optimizations. Secondly, it takes advantage of well-established frameworks like PyTorch. This enables the utilization of highly efficient parallel solvers, such as the GPU-accelerated Adam solver [23].

However, it is worth noting that the PINN approach may exhibit a trade-off between computational efficiency and accuracy compared to the CEGM approach [15]. It is hypothesized that the constraint-enforced nature of CEGM can lead to higher accuracy in estimating material properties. The physics informed loss terms of PINNs provide additional information to train the neural network, but do not inherently constrain it. Properly weighting the loss terms may improve accuracy, but is not be a generally applicable solution. Similarly, proofs on the convergence of CEGM require exact statically admissible stress fields, and while the PINNs approach uses the same constitutive error measure, it does not strictly ensure admissibility of these fields. Although some publications relax the static boundary condition, as they assume that measurement noise exists, they still strictly enforce internal equilibrium [20]. There are also additional explanations for the reduced accuracy. Aldirany et al. [24] show that it can also be explained by the inability of the NN optimizers to resolve the higher frequencies in the solution. This work enforces admissibility and constraints more strictly, as it hypothesizes that this is the main cause for the lack of accuracy.

Another drawback, not addressed in this publication, is the natural smoothness of NNs. The true constitutive and stress fields of a domain with material inhomogeneity are piecewise continuous (piecewise-C^0). Although the universal approximation theory [25] proves that NNs can fit such discontinuous jumps, they are not a particularly effective basis to describe such functions [26]. Alternative descriptions of these fields might require less degrees of freedom (DOFs) to fit the true fields with a similar accuracy.

3. Methodology and improvements

The primary objective of this study is to introduce a methodology that integrates the principles of CEGM into the PINN framework. This approach seeks to establish a formal PINN structure, wherein the constraints of CEGM are more rigorously enforced, with the hypothesis that this results in enhanced accuracy compared to previous PINN approaches.

3.1. Constrained optimization

Within CEGM, static admissibility is strictly enforced, emphasizing its importance over the error in the constitutive equation [20, 22]. Conversely, the PINN approach penalizes both aspects simultaneously, weighted with penalty factors. The reliance on penalty factors in penalization-based methods introduces challenges in hyperparameter selection. Excessive penalty factors can lead to spurious local minima and suboptimal solutions, while overly lenient factors can result in solutions that fail to satisfy the constraints. The application of the Lagrange multipliers method provides a solution to this challenge [27].

In the context of the inverse problem, the formulation is expressed as follows[1]:

$$\begin{aligned} \underline{\tilde{u}}^\star, \underline{\underline{\tilde{\sigma}}}^\star, \underline{\underline{\underline{\underline{\tilde{A}}}}}^\star = & \underset{\underline{\tilde{u}} \in \mathcal{N}^u, \underline{\underline{\tilde{\sigma}}} \in \mathcal{N}^\sigma, \underline{\underline{\underline{\underline{\tilde{A}}}}} \in \mathcal{N}^A}{\operatorname{argmin}} \mathcal{L}_{\text{cegm}} + \alpha \mathcal{L}_{\text{data}} \\ \text{s.t.:} \quad & \mathcal{L}_{\text{eq}} = 0 \\ & \mathcal{L}_{\text{bc}} = 0, \end{aligned} \tag{14}$$

which is solved using the Lagrange multipliers method, and seeks to find the stationary points of:

$$\mathcal{L} = \mathcal{L}_{\text{cegm}} + \alpha \mathcal{L}_{\text{data}} + \lambda_{\text{eq}} \mathcal{L}_{\text{eq}} + \lambda_{bc} \mathcal{L}_{\text{bc}}. \tag{15}$$

Although this looks similar to the penalization approach above, the difference is that multipliers λ_{eq} and λ_{bc} are additional DOFs in the optimization and will be determined automatically. This reduces the need for manual tuning of penalty factors and mitigates issues arising from magnitude differences between stress and displacement fields, common in inverse problems. The penalization parameter α is intentionally retained, as it can be used to train based on noisy full-field measurements. Extensive literature exists regarding the selection or even automated tuning of this parameter in traditional CEGM methodologies [20–22].

The constraints are applied to the internal and boundary equilibrium loss equations rather than formulating a separate constraint for each physics point for computational efficiency purposes. Adding a constraint for each point results in many Lagrange multipliers to be optimized for, and the optimization framework is not so well-equipped to handle this many constraints efficiently.

It is crucial to clarify that this approach should not be confused with the rigorous constraints inherent in the CEGM methodology. In CEGM, static admissibility is enforced across the entire stress field, whereas in this framework, static admissibility is numerically enforced exclusively at the points $\underline{x}_j$, where the loss functions are evaluated. Nevertheless, as the network is constrained at enough points, it is expected to uphold constraints throughout the entire domain. This assertion can be validated after optimization by assessing the constraints at sample points not considered during the optimization process.

3.2. Enforcing physics onto PINNs

More stringent formulations carefully design the PINN layout to exactly satisfy certain requirements. The following explains how thermodynamic admissibility and internal equilibrium equations can be satisfied precisely, without informing the neural network about these physics, and without requiring any training.

---

[1] A kinematic boundary condition can also be introduced, if necessary due to noisy or unreliable full-field measurements.

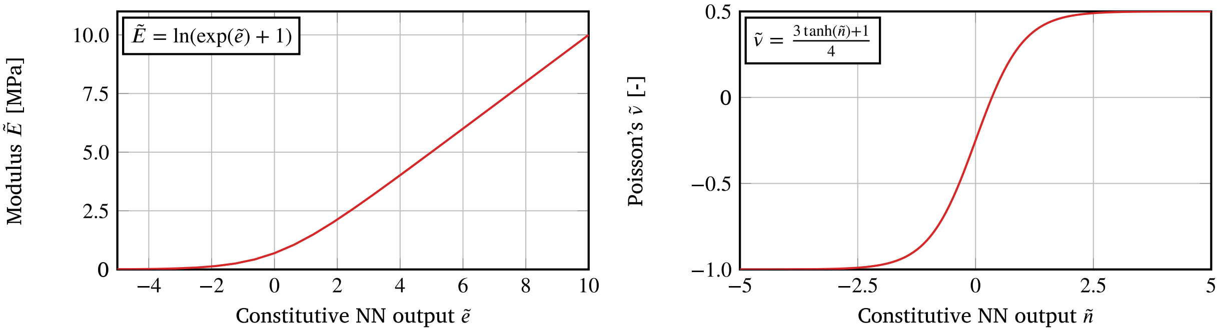

Fig. 3. Limiting the outputs of the constitutive network to an admissible domain. For these examples, the initial guesses E_0 and ν_0 were zero.

3.2.1. *Thermodynamic admissibility of the stiffness tensor*

The first improvement involves refining the formulation of the stiffness tensor. Wei et al. [15] assumed Hooke's law, describing the stiffness tensor with Young's modulus $\tilde{E}$ and Poisson's ratio $\tilde{\nu}$, scalar outputs of the constitutive NN. The NN allowed these variables to take any value, potentially resulting in stiffness tensors that violate the constitutive admissibility requirements. The constitutive equation gap, however, relies on the properties of an admissible stiffness tensor to construct a positive definite loss function. Thus, the method of Wei et al. [15] did not always converge. Depending on the seed (initial guess) of the NNs it might diverge, which is attributed to a negative CEGM when the stiffness is negative.

To address potential issues with convergence, the Young's modulus and Poisson's ratio are bounded to certain domains, ensuring positive definite stiffness tensors [28]. Instead of simple capping, a smooth function is applied to ensure admissibility while maintaining differentiability, as illustrated in Fig. 3. The

- **Young's modulus ($\tilde{E}$):** is required to be positive. Thus, the neural network output related to stiffness $\tilde{e}$ undergoes the transformation

  $$\tilde{E} = \ln(\exp(\tilde{e} + e_0) + 1), \tag{16}$$

  to find the actual Young's modulus. This function behaves linearly for large positive numbers, but has a lower bound of zero $\lim_{\tilde{e}\to-\infty} = 0$, as shown in Fig. 3. The initial guess E_0 can be used to reduce computational costs through $e_0 = \ln(\exp(E_0) - 1)$.
- **Poisson's ratio** ($\tilde{\nu}$): should never exceed 0.5 nor fall below −1, achieved through

  $$\tilde{\nu} = \frac{3}{4}\big(\tanh(\tilde{n} + n_0) + 1\big), \tag{17}$$

  where $\tilde{n}$ is the neural network output. Here, the initial guess (ν_0) is used to set $n_0 = \text{arctanh}(2\nu_0)$.

The equations above allow for the introduction of initial guesses. PyTorch uses Xavier Uniform Initialization [29] when seeding the weights of tanh-based fully-connected NNs, and sets the bias values to zero. Provided with small inputs, around unity.[2] This results into roughly zero-mean outputs with low variance and an approximately linear response. It is computationally more efficient when the true values are closer to this initial guess and thus this somewhat linear regime.

The initial guesses of this work were all homogeneous, and simply corresponded to the average stiffness of the sample. In a uni-axial tension example, this average stiffness was calculated from the specimen size, bench displacement and the force of the load cell, $E_0 = \frac{F}{A}\frac{L}{\Delta L}$ Alternatively, data normalization could be considered, see Appendix for a formal description. As it scales all fields towards unity, the average stiffness, and thus the initial guess, is 1.

The flat regions in the Young's modulus and Poisson's ratio limiting functions could cause the optimization to get stuck as the gradients vanish, but this is unlikely to happen. The problem of vanishing gradients is quite common in NNs with activation functions such as tanh or softplus. This potential ill-posedness is mitigated through normalization and initialization around zero. Although the flat region of the Young's modulus limit function is unlikely to be explored by the optimizer, applying this methodology to samples with (near) incompressible behaviour can cause issues. In those cases, the real Poisson's ratio is 0.5, and due to the limiting function, this might be approached but will not be reached. When dealing with incompressible materials, additional regularization or another parametrization of the constitutive equations should be considered.

For other constitutive models or parametrizations of the linear elastic model, a similar structure can be followed, determining the range of admissible constitutive parameters and choosing a function to ensure their properties remain within that range. Although not necessary, selecting functions that behave somewhat linearly within the expected value range may benefit optimization.

3.2.2. *Internal equilibrium*

Another violation of the admissibility requirements that can be alleviated is internal equilibrium. Instead of teaching the static neural network that the stress field should be divergence-free, it can be enforced through potential equations. The Airy potential

---

[2] Normalization and centring is required when the fields, such as the coordinates $\underline{x}$, become large, see Appendix.

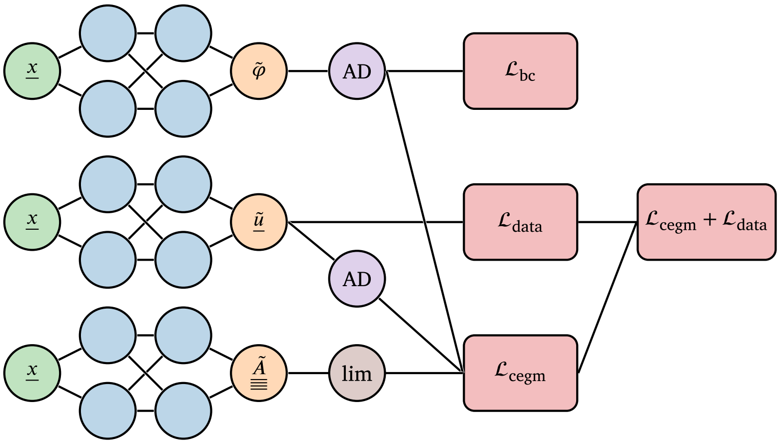

Fig. 4. Stress potentials such as the Airy stress for two-dimensional problems can be employed to ensure internal equilibrium. The actual stresses are computed from the potential using automatic differentiation (AD), while the constitutive properties are limited to a specific range (lim node) to ensure thermodynamic admissibility of the stiffness tensor.

method provides a systematic approach to formulating symmetric, divergence-free, second-order tensor fields in two dimensions using a scalar field [30]. This scalar field, $\tilde{\varphi}$, can be represented by a NN, from which stresses are computed using AD, as is proposed by Fig. 4:

$$\tilde{\sigma}_{xx} = \frac{\partial^2 \tilde{\varphi}}{\partial y^2}, \quad \tilde{\sigma}_{yy} = \frac{\partial^2 \tilde{\varphi}}{\partial x^2} \quad \text{and} \quad \tilde{\sigma}_{xy} = \tilde{\sigma}_{yx} = -\frac{\partial^2 \tilde{\varphi}}{\partial x \partial y}. \tag{18}$$

It can indeed be confirmed that the resulting stress field automatically satisfies the internal equilibrium equations everywhere, as it always is divergence-free and symmetric.

The requirement for the Cauchy stress field to be divergence-free arises from purely static considerations, independent of constitutive models or kinematic assumptions [31]. This principle holds even beyond the assumptions of infinitesimal strain and linear elasticity. In large deformation frameworks, the Cauchy stress remains divergence-free and symmetric [31], although this does not necessarily apply to stress measures in finite-strain Lagrangian formulations, such as Piola–Kirchhoff stresses.

A key challenge in ensuring equilibrated stress fields is handling potential material discontinuities. At these interfaces, the normal stress component tangential to the interface exhibits a jump, while the other components remain continuous. This discontinuity makes the divergence undefined at the interface, creating a singularity. It poses problems for the Airy stress potential, as well as for other PINN methods. Since these approaches create smooth, and enforce/train divergence-free stress fields, they struggle to capture the true, discontinuous nature of stresses at material interfaces. Nevertheless, Airy potentials can still be applied to problems involving material discontinuities. A common approach is to define separate Airy potentials for each subdomain, ensuring compatibility via interface traction conditions [32]. Equivalently, one may consider a single, piecewise-C^2 Airy potential, where discontinuities align only with the tangential normal stress component at the interfaces.

For inverse methods, the locations of these interfaces are generally unknown a priori, complicating the formulation. Two primary strategies exist: (1) defining internal boundaries along a predefined mesh or (2) approximating discontinuities weakly using functions that allow for steep gradients. The classical CEGM method [11] exemplifies the first approach, assuming homogeneous properties within each element and allowing stress discontinuities on the element interfaces while satisfying equilibrium. However, inaccuracies arise when the actual material interface does not align with the mesh. The PINN approach, by contrast, avoids a priori assumptions about material interfaces, leveraging the universal approximation theorem [25] to approximate even piecewise-C^0 fields. The Airy stress PINN similarly relies on the ability of NN to approximate piecewise-C^2 fields. Since the potential field is somewhat smoother, it is more accommodating to approximation with NNs.

Compared to traditional PINNs, Airy-PINNs provide two key benefits: they reduce the difficulty of fitting discontinuities at material interfaces while also ensuring internal equilibrium naturally, eliminating the need to explicitly train for this condition. However, the resulting stress fields still tend to be overly smooth, which remains the primary drawback of the proposed method, as it conflicts with fundamental principles of solid mechanics.

There are two more minor drawbacks related to computational costs: the need for the computation of expensive second derivatives and the larger neural network required to account for disregarded linear DOFs. Disregarding the linear DOFs leads to non-unique solutions of the potential field, and an ill-posed optimization problem as certain degrees of freedom have zero sensitivity. However, this does not cause convergence problems in practice. The Adam optimizer updates parameters with a step proportional to the exponential moving averages of the gradient and squared gradient [23]. When evaluated with AD, these are zero for every iteration, and the optimizer will take a step of size zero. Essentially, these parameters are “frozen” to the initial guess throughout the entire optimization process, avoiding any convergence problems. While the potential function itself may not be unique and depend on the initial guess, the resulting stress field is nonetheless unique.

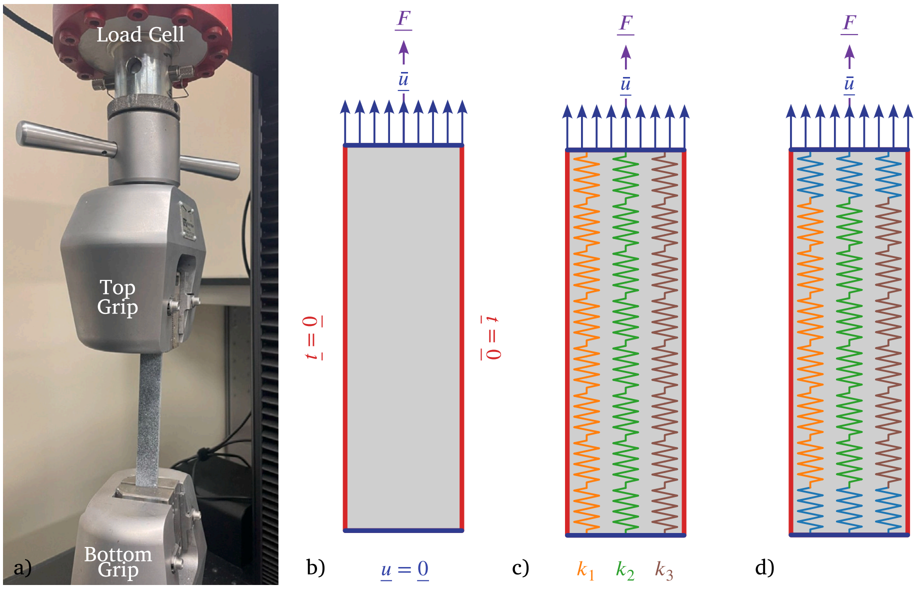

Fig. 5. A tensile test with a universal testing machine in (a) photo and (b) schematic. The sample is gripped on kinematic boundaries over which the total reaction force is measured with a load cell, while the other boundaries are traction-free. (c) shows how the problem becomes non-unique if the microstructure is perfectly parallel, whereas (d) recovers uniqueness by assuming constant constitutive properties along the kinematic boundaries.

3.3. *Solving the degenerate examples through regularization*

In this paragraph, we highlight some important consideration for cases without non-zero static boundary conditions, or when these are known only partially through integral forms. In the first case, the “scale” of the material stiffness cannot be determined, while the second case introduces additional complications. One typical example would be a simple tensile test on a perfectly aligned microstructure, as shown in Fig. 5. This cannot be solved with the procedures discussed above as it is a degenerate problem without a unique solution. Such experimental setups, that contain kinematic boundaries (for example. in the grips of the load frame across which both the displacement and total force are measured, while all other boundaries have zero traction prescribed) require special treatment as defined below.

The formulation can be extended to include such boundary conditions [20]. This requires the introduction of two constraints. Firstly, the discrepancy between the NN displacement and prescribed displacement boundary condition. Secondly, a static condition that integrates the traction along the boundary and compares it to the reaction force $\underline{F}$ measured with the load cell:

$$\mathcal{L}_{bc_1} = \sum_{k=1}^{N_b} \left\| \underline{\tilde{u}}(\underline{x}_k) - \underline{\hat{u}}(\underline{x}_k) \right\|^2 \Delta S_k \, , \quad \text{and} \quad \mathcal{L}_{bc_2} = \left\| \sum_{k=1}^{N_b} \underline{\underline{\tilde{\sigma}}}(\underline{x}_k) \cdot \underline{n}(\underline{x}_k) \Delta S_k - \underline{F} \right\|^2 , \tag{19}$$

where $\underline{x}_k$ are points on this boundary. These equations are added as additional constraints to the Lagrange multipliers formulation introduced in Section 3.1.

Some problems remain difficult to solve, even with these additions. Fig. 5 illustrates an example in which this formulation clearly lacks a unique solution. Due to the homogeneous kinematic constraints, all springs deform equally, making it impossible to determine which has the highest or lowest stiffness. If the local traction distribution were known, the problem could be solved, as the stiffest springs carry the highest load.

Sections 4.3 and 5 include examples of tension test results where only the total reaction force is available, rather than local traction distributions. Although there is not a complete loss in uniqueness, the lack of local traction distribution information will make these problems difficult to solve. To address this issue, two regularization methods are introduced to reformulate the problem: the first ensures uniqueness under all circumstances, while the second, a simpler method, is effective only under specific conditions.

Both proposed solutions are not just limited to the PINN-based formulations discussed in this paper, but can also be implemented into traditional CEGM approaches to increase the accuracy of inferred traction information and subsequent material identification. The development of a method solving the uniqueness problem was required for the real experimental examples discussed in Section 5, as the PINN procedure failed to solve these problems without additional regularization.

3.3.1. *Multiple loading conditions method*

This generic approach adds more information — specifically, the deformation of independent loading conditions, as illustrated in Fig. 6. For a two-dimensional example, combining tensile loads in directions $\underline{e}_y$, $\underline{e}_x$, and a shear test provides multiple sets of displacement measurements and static boundary conditions and enough information to recover the constitutive property distribution.

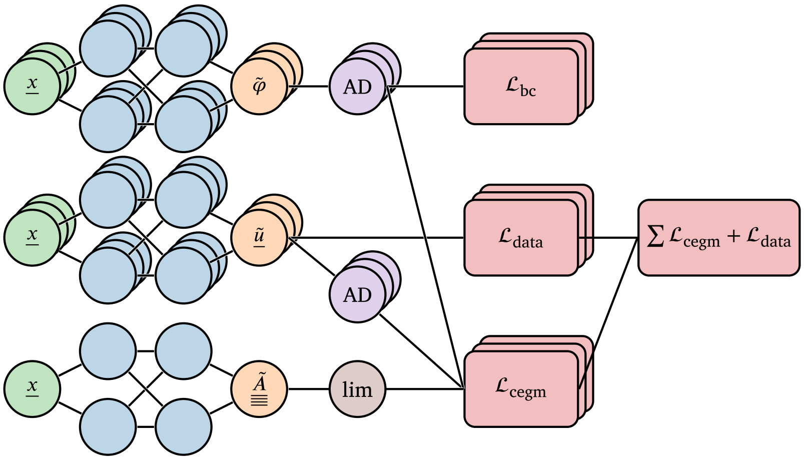

Fig. 6. Solving multiple independent load cases requires several kinematic and static NNs, but only a single constitutive one, assuming that the coordinate system $\underline{x}$ of the different measurements matches with one another.

This is akin to a homogenization problem, which similarly requires different loading configurations to determine the overall stiffness tensor [33].

This approach is generic but difficult and expensive to apply. Firstly, it is experimentally costly as it requires equipment to apply these different loading conditions onto the sample. One also has to ensure that the coordinate system matches between the different experiments (possibly through marks on the specimen). Lastly, the inverse PINN concerns many more NNs and DOFs, making it more expensive to optimize for.

3.3.2. Homogeneous regions method

In cases where the specimen is isotropic and homogeneous along the kinematic and total force boundary conditions, a more efficient approach can be employed. This method regularizes the problem by explicitly adding the assumption of unknown yet homogeneous properties along the boundary, as illustrated in Fig. 5.

The PINN automatically deduces the traction distribution on the kinematic boundary from the measured deformation along the boundary. After all, constant stiffness along the boundary means that the region with the most deformation was subjected to the highest traction.

The implementation, as depicted in Fig. 4, undergoes minimal modification. When the elastic properties for any point $\underline{x}$ in the homogeneous region are required, parameters from a vector are provided rather than the values of the constitutive neural network at $\underline{x}$. These parameters are then introduced as degrees of freedom during the training process. This method stands as the preferred approach when applicable, as the subsequent problem solves quickly, while maintaining low experimental costs as it requires only a single test.

4. Numerical results

This section presents a comparative analysis of in silico examples of the proposed methodology in conjunction with two established approaches: the CEGM method [11] and the PINN method [15] that we seek to improve. In contrast to previous work, it will include examples with kinematic boundary conditions, which were difficult to solve due to the absence of local non-zero traction information.

4.1. Example problems and network design

For the purpose of evaluation, reference problems from prior research are employed, as depicted in Fig. 7. No additional noise was introduced in these numerical examples. These problems involve non-homogeneous 10 by 10 mm squares, and the displacement field is computed using Abaqus on a highly refined mesh, subsequently down-sampled to 11 points in each direction. This coarse resolution is inherited from the CEGM examples due to computational limitations [11]. Despite the displacement measurement resolution of 1 mm, the constitutive equation gap loss and static boundary constraints can be assessed at any location. A spatial resolution of 0.33 mm (31×31 points) was determined to provide accurate results without overfitting. This resolution is increased compared to the reference study, as that publication had minor overfitting problems.

To extract the most information possible from the data, one would like the functional space described by the NNs to be as large as possible. This is done by increasing the number of layers and nodes per layer, and essentially allows the description of more refined displacement, stress and constitutive fields. However, NNs that are too large will result in overfitting. This occurs when a model

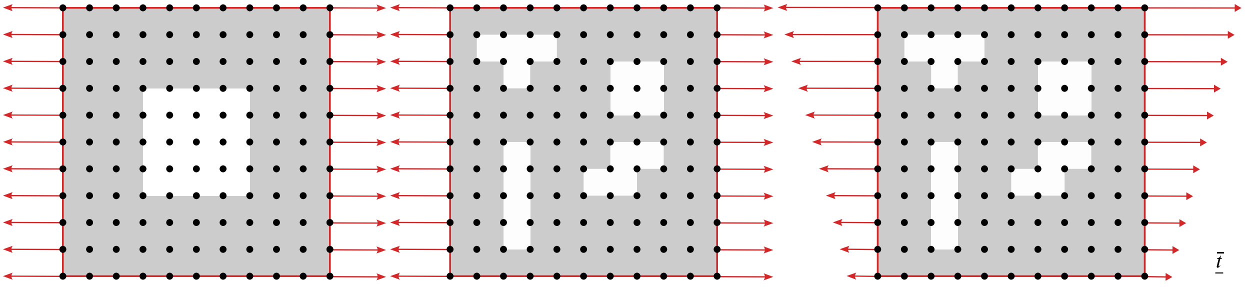

a) RP1	b) RP2	c) RP3

Fig. 7. The three reference problems, where the black dots represent measurement points. The white regions have a stiffness of 2 MPa, whereas the grey areas are 1 MPa. The Poisson's ratio is constant at $\nu = 0.3$ throughout the entire domain.

contains too many degrees of freedom compared to the available data, resulting in solutions that fit the training data, including its noise or outliers, excessively well, but generalize poorly to unseen data.

This overfitting should be prevented as the results become non-physical. The following procedure was followed to select the NN-size and amount of data such that overfitting did not occur:

1. **Initial choice** of NN-size (number of layers and nodes per layer) and data-size (number of physics-points and displacement data-points from DIC).
2. **Train** the NNs until all losses converge.
3. **Test** all losses on previously unseen locations (other physics-points and new displacement-points):
   (a) **Overfitting is detected if any losses increase significantly.** Either increase amount of data (physics and/or displacement) or decrease NN-size and start at point 2 again.
   (b) **No overfitting is detected if losses remain unchanged.** Therefore, results can be relied upon.

For consistency with prior work, the NN architecture closely follows that of Wei et al. [15]. Consequently, all networks are fully connected and employ a hyperbolic tangent (tanh) activation function. The implementation of NNs is achieved through PyTorch [34], with optimization carried out using the MDMM algorithm for stationary point optimization[3] [27]. The training initializes with a learning rate of 0.01, with automatic reduction upon reaching plateaus, while the optimization process concludes after 40,000 epochs. The PINN results [15] are recomputed using the same settings to ensure a fair comparison.

The displacement, Cauchy stress, and constitutive networks possess three hidden layers with 32 nodes per layer, while the Airy stress potential network comprises four layers with 64 nodes per layer. Such a large network for Cauchy stress would result in overfitting, whereas this does not occur for the Airy stress formulation. This is because the Airy stress network results in less freedom than a Cauchy stress network of the same size, as only the second derivatives of the network are used, removing the linear contributions.

The accuracy is visualized using the relative error,

$$e^{\square}(\underline{x}) = \frac{\left|\tilde{\square}(\underline{x}) - \square^{\mathrm{ex}}(\underline{x})\right|}{\left|\square^{\mathrm{ex}}(\underline{x})\right|}, \tag{20}$$

which compares the NN approximation to the true solution at any location $\underline{x}$, resulting in the heatmaps shown below. The $\square$ represents the variable of interest, either Young's modulus E or Poisson's ratio ν. To aid comparison between different methods, the mean and standard deviation of this error field are obtained through integration over the space. The mean represents the average relative error over the entire domain, while the standard deviation provides information on the uniformity of the error.

$$e_{\square}^{\mathrm{mean}} = \frac{1}{\mathrm{vol}(\Omega)} \int_{\Omega} e_{\square}(\underline{x})\, \mathrm{d}\underline{x} \quad \text{and} \quad e_{\square}^{\mathrm{std}} = \sqrt{\frac{1}{\mathrm{vol}(\Omega)} \int_{\Omega} \left| e_{\square}(\underline{x}) - e_{\square}^{\mathrm{mean}} \right|^{2}}. \tag{21}$$

4.2. Constrained optimization and enforced physics

To assess whether each change offers any advantages compared to existing literature, two variants of the algorithm are created, and their results are listed in Table 1. Both implementations contain the thermodynamically admissible stiffness tensor. The first variant, hereafter referred to as "Constrained PINN", directly utilizes a stress field obtained from a NN. It employs constrained optimization with Lagrange multipliers and enforces the thermodynamic admissibility of the constitutive equations. The second variant, dubbed "Airy PINN", employs the stress potential formulation, and uses the Lagrange multiplier only for the boundary equilibrium loss.

---

[3] The Lagrange multiplier optimization method for automatic differentiation frameworks (MDMM) is not implemented in Torch, but "crowsonkb" has made a small library (https://github.com/crowsonkb/mdmm) that wraps the saddle-point optimizer required for MDMM in Torch.

Table 1
Comparing the mean and standard deviation of the relative error among the different methods from measurements with 1 mm resolution.

	RP1		RP2		RP3		Time
	$\langle e_E \rangle$ [%]	$\langle e_\nu \rangle$ [%]	$\langle e_E \rangle$ [%]	$\langle e_\nu \rangle$ [%]	$\langle e_E \rangle$ [%]	$\langle e_\nu \rangle$ [%]	t [h:mm]
CEGM	2 ± 7	4 ± 5	5 ± 8	7 ± 7	5 ± 8	7 ± 7	>4:00
Traditional PINN	8 ± 9	8 ± 8	13 ± 13	10 ± 10	14 ± 13	10 ± 11	0:05
Constrained PINN	6 ± 8	6 ± 6	10 ± 13	11 ± 11	11 ± 13	12 ± 11	0:07
Airy PINN	5 ± 8	5 ± 6	9 ± 12	9 ± 8	9 ± 12	9 ± 8	0:15

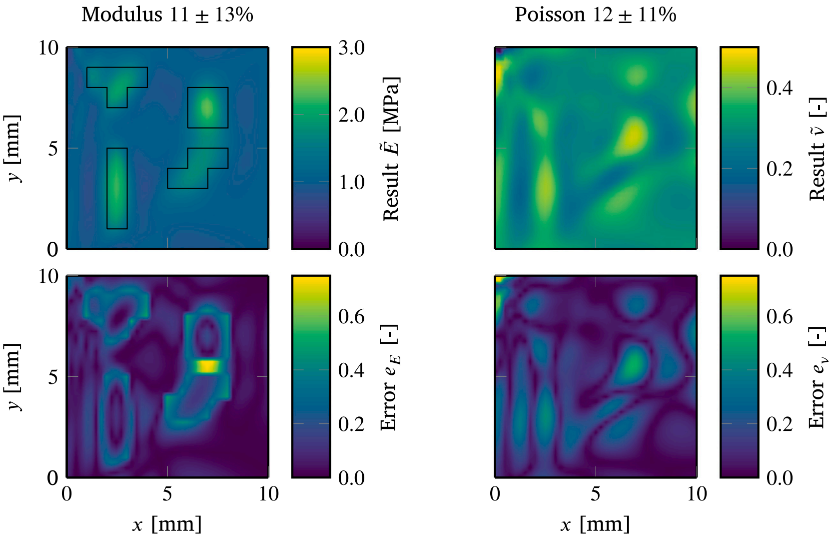

Fig. 8. Constitutive properties and errors of RP3 obtained with the constrained PINN.

No example independently shows the impact of thermodynamically admissible stiffness tensors, as that does not impact the accuracy of the results themselves. The thermodynamically admissible stiffness tensor formulation is, however, much more reliable in terms of convergence. The traditional PINN, either converges to positive or diverges to negative stiffness depending on the NN seed. No divergence was observed for any seed after the constitutive properties were limited to an admissible domain, although the accuracy of the solution $\langle e_\square \rangle$ still varies by ±1%, depending on the selected seed. To ensure repeatability and consistency, all simulations reported initialize from the same seed, selected through trial and error such that the traditional PINN also converges for all examples.

The overview in Table 1 indicates that the constrained formulation slightly improves the PINN method in most cases, but not all. Whereas the Airy stress formulation consistently outperforms the other PINNs, though it does not reach the accuracy of the traditional CEGM method. The resulting constitutive distributions of RP3 are visualized in Figs. 8 and 9, with the true stiff inclusions indicated. The stiffness obtained through the coarse Airy stress approach, as depicted in Fig. 9, effectively separates the stiffer inclusions. Conversely, the constrained formulation, as seen in Fig. 8, struggles to differentiate the two inclusions on the right side of the domain. All PINN approaches, even the traditional PINN which is not plotted, suffer from overly smooth constitutive property fields. The recovery of stress fields through these inverse methods is less accurate than the recovery of material properties. The stress fields showcased in Fig. 10 merely capture the overall trends of the stress fields, while intricate details and discontinuities are smoothed out by the neural network. Although the Airy based stress field is still too smooth, it is much closer to the exact solution. This is because approximating the piecewise-C^2 potential field is easier with NNs than the piecewise-C^0 stress field used in the constrained and traditional PINN.

A deeper investigation reveals that the NNs are simply not rich enough to capture the exact constitutive and static fields, especially around the discontinuities at the material interfaces. Directly training the constitutive neural network on an exact and refined distribution of the constitutive properties results in a Young's modulus approximation with an accuracy of 2.5%. Similarly, training the static NN directly on the exact stress fields results in a mean relative error of 3%. It was concluded that the constitutive and static neural networks cannot accurately represent the discreteness of the true fields, because (small) NN are inherently smooth. Regardless of the solving strategy, the neural network representation does not have enough freedom to predict the constitutive properties with the accuracy of the CEGM method.

To address this, the adoption of larger neural networks might yield better results. A NN, with three layers and 48 nodes each, can fit the exact stiffness distribution to a 0.8% relative mean error, while it could fit the exact Cauchy stress to < 1%. An Airy potential NN needs to be larger, at least 92 nodes, for that accuracy. However, increasing the network size results in overfitting

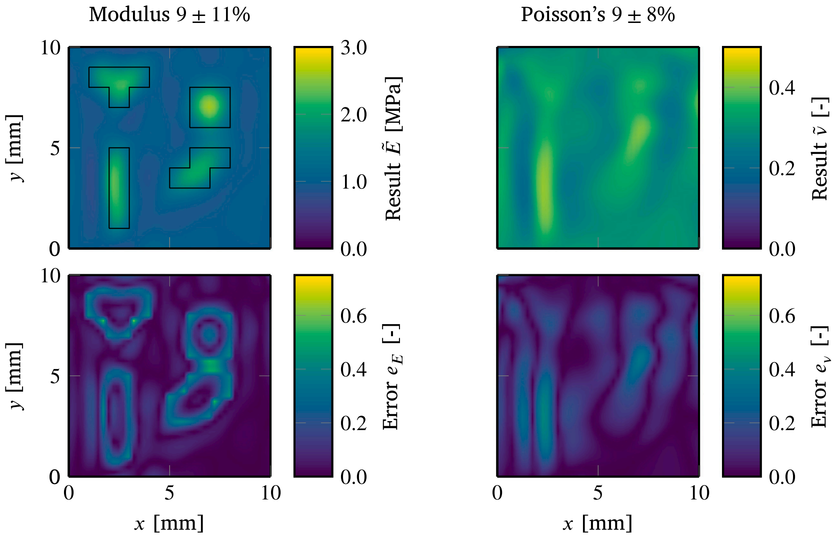

Fig. 9. Constitutive properties and errors of RP3 obtained with the Airy potential PINN.

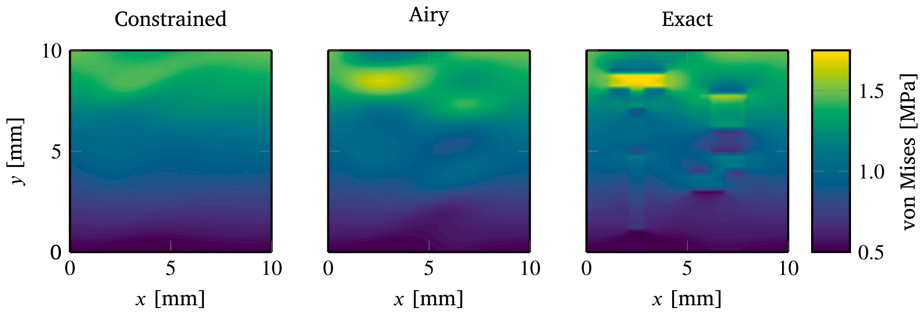

Fig. 10. Comparing the von Mises stresses obtained through the different approaches using the coarse simulations.

Table 2
Comparing the mean and standard deviation of the relative error among the different methods from displacement fields with a 0.25 mm resolution.

	RP1		RP2		RP3		Time
	$\langle e_E \rangle$ [%]	$\langle e_\nu \rangle$ [%]	$\langle e_E \rangle$ [%]	$\langle e_\nu \rangle$ [%]	$\langle e_E \rangle$ [%]	$\langle e_\nu \rangle$ [%]	t [h:mm]
Traditional PINN	7 ± 8	7 ± 8	10 ± 12	8 ± 7	13 ± 13	7 ± 7	0:11
Constrained PINN	4 ± 8	5 ± 5	10 ± 12	10 ± 10	8 ± 11	9 ± 8	0:18
Airy PINN	3 ± 7	3 ± 4	6 ± 10	5 ± 6	7 ± 10	5 ± 5	0:45

problems, and adding more physics points does not resolve the situation; thus more measurement samples are required. A finer displacement measurement resolution is acceptable, as the 1 mm resolution was inherited from the original CEGM paper [11] and does not reflect the capabilities of modern DIC systems. In a study by Latourte et al. [35], a simple 8 MP camera was used, already resulting in displacement fields with a much finer resolution. Similarly, the stereo-DIC examples in this publication were obtained with two 12.3 MP cameras. When projected onto the specimen, the resolution is approximately 0.02 mm.

A summary of the results for simulations with larger NNs and the refined measurement grid are listed in Table 2. The resolution of the displacement field is reduced to 0.25 mm (41×41 points), while the physics losses are sampled by 61×61 points. Here, the accuracy approaches, and for the Airy potential formulation even exceeds, what is obtained with the CEGM method on the coarse grid. This shows that larger NNs become more capable of capturing the required discreteness.

Tables 1 and 2 also provide summaries of the average computational costs.[4] All PINN approaches outperform the CEGM method, although the Airy-based formulation is more expensive than the others. The increased computational cost associated with the

---

[4] All simulations were performed on a Lenovo Yoga Pro 7 laptop with an AMD R7 7735HS processor and 16 GB RAM and a Fedora 40 operating system All GPU acceleration was disabled.

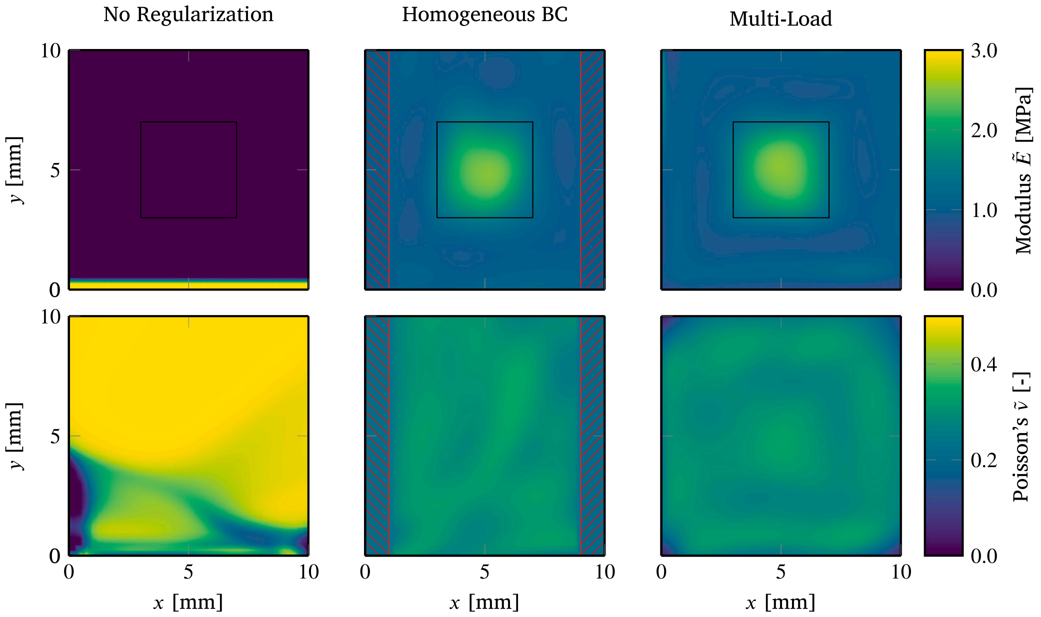

Fig. 11. The Young's modulus and Poisson's ratio of the example problem with a total force instead of local traction on all non-zero traction boundaries. The regions with assumed homogeneous properties along the boundary are hashed, while the stiffer region in the middle is annotated in the results of all methods.

Airy stress formulation is attributed to two primary factors. Firstly, it necessitates the computation of higher derivatives, leading to a more intricate chain of automatic differentiation. Secondly, larger networks are required to ensure adequate flexibility for accurately describing the stress field, as all linear components of the network are disregarded in such a potential formulation, introducing more unknowns. The constrained formulation on the finer mesh has a similar performance, both in accuracy and speed, as the Airy potential method had on the lower resolution dataset. It is important to emphasize that these additional computational costs of the Airy potential method translate into improved predictive accuracy, and without the need for larger datasets. For the best accuracy, the Airy stress potential method is recommended, especially in cases with limited resolution in the full-field measurements. For minimal computational costs, with high resolution data available, the methods are equivalent if the Airy potential method is performed on a down-sampled dataset. Lastly, the iterations involving constraints tend to be slightly more resource intensive. Consequently, in scenarios where the number of iterations are fixed, constrained optimization is expected to incur higher computational costs, explaining why the PINN method in the literature is cheapest. However, the increased costs are insignificant when contrasted with the computational burden of the traditional CEGM method.

4.3. Solving the typical tensile setup without local traction information

Three inverse problems were solved to illustrate the proposed remedies for the uniqueness problem stemming from the lack of detailed traction information. These setups have traction distributed over a boundary of which only the total force is known. The first strategy naively implements the kinematic and total force boundary conditions as a Lagrange multiplier constraint, see Eq. (19). This shows how poor the results are when the algorithm is not assisted with inferring the local traction boundary condition. Then the constraint of homogeneous properties along the boundary conditions are introduced, resulting in accurate predictions. This is followed by the multi-load case method, which simultaneously solves for the three independent measurements.

The considered scenario concerns identical material distribution as RP1 in Fig. 7, but now with loading through kinematic conditions. For the first two examples, a kinematic boundary was imposed such that the left edge was clamped, while the right side was displaced by 1 mm in $\underline{e}_x$ direction, and a reaction force $\underline{F} = 11.5\underline{e}_x$ N was obtained from FEA nodal-force results along that boundary. For the multi-load case method, this was augmented with two additional tests. A tensile test where the bottom was clamped, and the top boundary was displaced by 1 mm in $\underline{e}_y$ direction, resulting in a load $\underline{F} = 11.5\underline{e}_y$ N. Subsequently, the specimen was loaded in shear, where the left boundary was clamped, and the right boundary displaced 1 mm upwards, causing a $\underline{F} = 2.9\underline{e}_y$ N load to be observed.

The PINN architecture use the Airy stress formulation, while the network size follows the description in Section 4.1. Some changes needed to be made to the optimization settings, since these problems proved more difficult to solve. Although, the displacement is still measured on an 11×11 points grid, the number of points where the CEGM equation is sampled needed to be increased to 61×61 points to prevent overfitting. Similarly, the number of iterations needed to be increased to 75,000 epochs, as especially the multi-load case method converges slower. For consistency, these settings were used for all these inverse problems. This also results in slightly higher computational costs, as acknowledged in Table 3.

Table 3

The accuracy of solutions that have kinematic boundary conditions on which the only non-zero static component is the reaction force.

	$\langle e_E \rangle$ [%]	$\langle e_\nu \rangle$ [%]	t [h:mm]
No regularization	192 ± 545	55 ± 19	0:45
Homogeneous boundary	8 ± 9	10 ± 7	0:47
Multi-load case	8 ± 9	6 ± 9	2:23

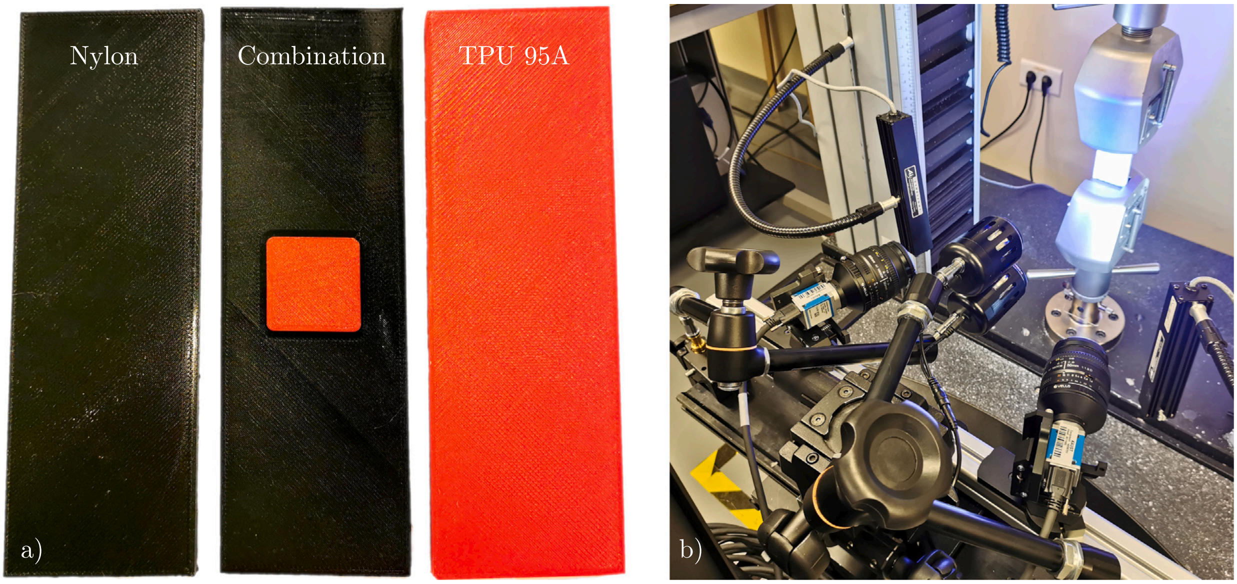

Fig. 12. (a) From left to right, the homogeneous nylon, inhomogeneous, and homogeneous TPU 95A samples, and (b) the test setup.

Table 3 also compares the performance of the different implementations. Clearly, the implementation that only adds the appropriate boundary condition does not result in the correct solution. It is also not unique, as the result depends strongly on the initial guess. However, just assuming homogeneous properties along the clamped boundaries is enough to converge to the correct solution. The much more expensive, multi-load case approach does not predict the Young's modulus more accurately, but increases both the experimental complexity and computation costs.

Note that the Poisson's ratio is less accurate than in previous examples, especially for the homogeneous boundary method. This discrepancy is caused by the clamped boundary. In the centre region, everything except the hashed areas in Fig. 11, the Poisson's ratio is much more accurate, resulting in $\langle e_\nu \rangle = 5 \pm 5$, which is comparable to previous findings. At a clamped boundary, Poisson's effects cannot be observed, as the boundary condition restricts the kinematic freedom. Thus, the inverse methods under-predict the Poison's ratio. This is also why the Poisson's ratio of the multi-load case method is inaccurate at the corners, as those locations are dominated by the clamped boundaries in all load-cases considered.

It is concluded that both methods successfully identify the material property distributions. The homogeneous boundary method is preferred due to the lower costs, but can only be used if that assumption is realistic. Otherwise, the multi-load case approach is the only option. Lastly, it is recommended to exclude the Poisson's ratio results close to the clamped boundaries, as those results are less accurate.

5. Experimental example

Samples with inhomogeneous materials were manufactured and tested to demonstrate how the proposed method can identify material properties under realistic experimental conditions.

5.1. *Test methodology*

Three specimens were produced: one from nylon, one from TPU 95A, and one combining both materials, as shown in Fig. 12. All samples were produced using additive manufacturing with an Ultimaker S5 at 100% infill density, with all other settings at default. The samples had a thickness of 5 mm and a width of 50 mm. The inhomogeneous sample featured a 25 mm soft square inclusion with rounded corners of 2 mm.

To prepare for DIC, the samples were painted white, and black speckles were applied using an airbrush. DIC was used, see Fig. 12, to obtain the displacement field under loading. The stereo-DIC setup consisted of two Sony IMX253 cameras with Nikon 50 mm

1.8 D lenses. The tensile machine grips were 75 mm apart, with approximately 70 mm of the sample height visible through the cameras. A displacement-controlled loading rate of 2 mm/min was applied for a total displacement of 1.5 mm, corresponding to an average strain of 2%. The load was recorded for each sample and will later be used as a boundary condition for the inverse method. Displacement fields were computed using Correlated Solutions VIC-3D software, and only the initial and final images were used for the inverse analysis.

5.2. Data preparation

The data preparation process involved three steps: downsampling, centring, and normalization. Downsampling was necessary as the DIC algorithm computed the displacement field at 9 million coordinates. The resulting datafile could not be loaded in system memory, let alone be operated upon. Even without this limitation, PINNs should never use the full dataset for training, as one should test against overfitting on data unseen by the optimizer. The down-sampled data was used during the training phase, while overfitting was tested on a fine grid of 501×701 points.

Centring was performed for convenience in defining boundary conditions. The coordinate system from DIC was shifted such that the image centre aligned with the origin. Normalization is essential, as the tanh-based NN and optimizers struggle with fields of vastly differing magnitudes. Normalization is standard practise when working with NNs, as such the details are listed in Appendix.

5.3. Identifying homogeneous specimens

The homogeneous samples were identified with two methods. Firstly, with hand calculations that compare the strain values measured with DIC to the applied loading. Secondly, through the Airy stress inverse method introduced in this publication.

The hand calculations estimated the constitutive properties using the strain field in the last DIC image. These strain fields were not perfectly uniform, as lines at $\pm 45°$were visible. These were probably caused by slight inhomogeneities from the additive manufacturing process. These defects can visually be observed on the sample in Fig. 12. It was assumed that these inhomogeneities average out, so the Young's modulus and Poisson's ratio were defined as:

$$E = \frac{1}{N}\sum_{n}^{N} \frac{\sigma_{yy}}{\varepsilon_{yy}(\underline{x}_n)} \quad \text{and} \quad \nu = \frac{1}{N}\sum_{n}^{N} -\frac{\varepsilon_{xx}(\underline{x}_n)}{\varepsilon_{yy}(\underline{x}_n)}, \tag{22}$$

where the stress σ_{yy} is the measured load-cell force divided by the cross-sectional area, while $\underline{x}_n$ refers to all locations in the DIC dataset. The spatial standard deviation of the constitutive properties will also be reported to acknowledge the slight inhomogeneity from the manufacturing procedure.

The Airy stress method was employed for analysing the homogeneous samples, with the same architecture and hyperparameters used as in the numerical examples from Section 4.1. The displacement and constitutive networks had three hidden layers of 32 nodes each, while the Airy stress potential network comprised four layers with 64 nodes each. The downsampled displacement field contained 25×35 points, and the physics were evaluated on the same grid. These problems contain clamped boundaries lacking local traction information. Therefore, the inverse method might not converge to a solution. As discussed in Section 4.3, this was solved by assuming a region of homogeneous material properties along those boundaries.

The training process began with a learning rate of 0.01, which was automatically reduced when plateaus were reached. Optimization was performed for a total of 40,000 epochs.

5.4. Identifying inhomogeneous specimens

For the inhomogeneous sample, the inverse-PINN did not converge under the initial settings. More iterations and larger networks were required.

The kinematic and constitutive neural networks were increased to 72 nodes per layer, while the static NNs layers contained 72 and 128 nodes for the constrained and Airy formulation, respectively. The displacement field was sampled at 101×141 points, and the physics were evaluated at 96×128 points to reliably fit the larger networks.

To reduce computational expense, the kinematic NN was first trained directly on the measured displacement field for 5000 epochs. This was followed by optimization of the other NNs for 10,000 epochs, during which the kinematic NN was fixed. This approach initialized the neural network with a better initial guess. Subsequently, all weights and biases were optimized for with another 400,000 epochs.

After training, a threshold method was used to determine the location and size of the soft inclusion. Areas with normalized stiffness below a certain value were classified as the soft region, while the rest was considered the hard phase. The impact of the threshold value will be reflected on. The spatial mean and standard deviation of the constitutive properties were calculated for the two regions, as they represent the material properties of the two phases. A drawback of this method is that it includes information from the smooth transition at the interface. This has a major impact on the spatial standard deviation. An alternative approach involves using two threshold values: an upper threshold above which points are classified as part of the stiff phase, and a lower threshold below which points belong to the soft phase, while intermediate values are treated as interface-related areas and thus disregard. This approach reduces the impact of the non-physical smooth transition but was ultimately not reported, as it could lead to an unfair comparison.

Table 4
Material properties and standard deviations obtained through the various inverse methods. A normalized stiffness cut-off of 0.5 was used to separate nylon and TPU regions in the inhomogeneous sample.

	Nylon		TPU 95A	
	E [MPa]	ν [–]	E [MPa]	ν [–]
Homogeneous samples with hand-calculation	1536 ± 60	0.40 ± 0.02	58 ± 6	0.43 ± 0.02
Homogeneous samples with Airy PINN	1539 ± 37	0.41 ± 0.01	61 ± 6	0.44 ± 0.01
Inhomogeneous sample with Constrained PINN	1622 ± 406	0.37 ± 0.10	127 ± 53	0.24 ± 0.11
Inhomogeneous sample with Airy PINN	1676 ± 294	0.42 ± 0.05	82 ± 48	0.27 ± 0.09

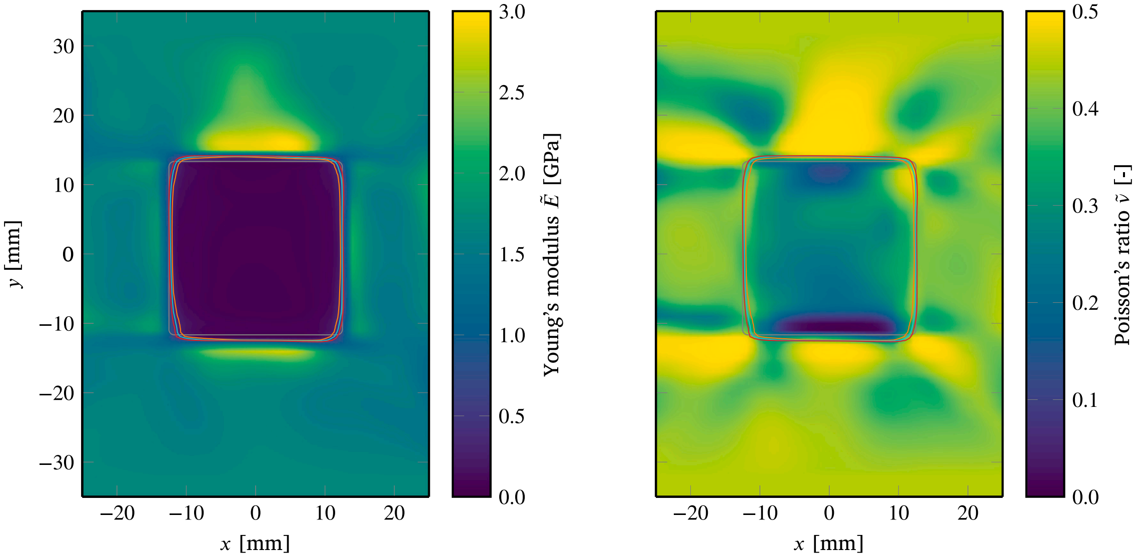

Fig. 13. Constitutive properties of the inhomogeneous specimen obtained with the Airy PINN method. The boxes represent the boundaries of the soft phase: the grey box shows the true boundary, while red, blue, and orange indicate boundaries corresponding to threshold values of 1.00, 0.75, and 0.50, respectively.

5.5. Results and discussion

The constitutive properties obtained for the homogeneous samples, listed Table 4, shows that the hand calculation and inverse-PINN closely matched the expected values for nylon and TPU. The Young's modulus for TPU was determined to be 61 ± 6 MPa, which is consistent with the 67 ± 6 MPa reported in the technical datasheet (TDS) [36]. The Poisson's ratio for both materials, 0.41 ± 0.01 for nylon and 0.44 ± 0.01 for TPU, aligns with the typical range for such polymers, but they are not listed on the TDS. The only outlier was the Young's modulus of the nylon, 1539 ± 37 MPa, which is lower than the supplier's reported 2331 ± 55 MPa [37], but consistent with other studies reporting values between 1630 and 1750 MPa [38]. This discrepancy could be due to different manufacturing parameters or moisture uptake in the nylon spool. A traditional Young's modulus extraction with simple hand calculations confirmed the findings obtained with the NNs-based inverse method.

It is also noteworthy that the inverse method correctly identified both specimens as homogeneous. Although the constitutive neural network allowed for non-uniform material properties, the constitutive fields obtained in both cases were nearly constant, as indicated by the small spatial standard deviations in Table 4.

Evaluating the inhomogeneous sample with the Airy PINN yields constitutive properties somewhat consistent with those of the homogeneous samples, as listed in Table 4, whereas the constrained PINN demonstrates lower accuracy. In fact, the results from the homogeneous samples fall within one standard deviation of those from the inhomogeneous sample, except for the Poisson's ratio in the TPU region. The lower Poisson's ratio observed in the TPU region is primarily due to underprediction at the bottom edge of the inclusion, as shown in Fig. 13. In reality, the Poisson's ratio field should be fairly uniform since both materials have similar Poisson's ratios. However, the recovered field is significantly non-homogeneous. This inaccuracy in Poisson's ratio predictions is not unusual, as it is typically more difficult to estimate than Young's modulus. Previous studies using CEGM, with [15] and without PINNs [11], encountered similar challenges. The in-silico examples of those studies also demonstrated inaccuracies in predicting the Poisson's ratio although only the Young's modulus field was inhomogeneous. These occur at the boundary between phases with different Young's moduli, which suggests that it might be the minor inaccuracies of the Young's modulus that trigger larger differences in the Poisson's ratio.

For both PINN-based methodologies, this is exacerbated by the non-physical smoothness in the constitutive and static fields at the material interface. Larger NNs might fit these discontinuities better, but this requires more displacement and physics points to prevent overfitting. This information is available (the downsampled dataset contains around 1% of all the points in the DIC field), but this becomes prohibitively expensive from a computational perspective. Especially memory limitations prevented the evaluation of larger models.

Because all fields were normalized in the data-preparation stage, an initial stiffness $E_0 = 1$ was used upon seeding the NN. As the stiffer Nylon phase dominates the global response, this initial guess was inaccurate for the soft TPU region as the true Young's modulus was a factor 25 lower. Nevertheless, the inverse-PINN correctly determined the magnitude of the TPU stiffness, showcasing that a bad initial guess does not affect the results.

The size of the soft inclusion, determined using a normalized stiffness threshold of 0.5, was found to be approximately 24 by 26 mm, which is close to the actual size of 25 mm square. Fig. 13 illustrates the boundaries of the soft inclusion for various threshold values. The boundaries are close to each other, which indicates a low sensitivity to the specific threshold value, suggesting that the selection of this hyperparameter is not critical. This low sensitivity can be attributed to the steep gradient in the stiffness field around the inclusion.

Overall, these results demonstrate that this type of inverse method can accurately extract material properties from real experimental tests. These experiments posed significant challenges, as the true behaviour of the specimens cannot be fully captured by the assumption of isotropic linear elasticity. The 3D-printed samples are inherently anisotropic, as they are constructed from layers with $\pm 45^\circ$ orientations and this was also visible in the strain fields obtained with DIC. Furthermore, the materials themselves, particularly TPU 95A, do not conform to linear elasticity. TPU is a soft plastic that typically requires more advanced, non-linear constitutive models. Future work could explore performing the same inverse analysis with a hyper-elastic material model, which may lead to more accurate results.

6. Conclusion

This study shows that the PINN approach to material identification can be improved by enforcing physics rather than merely informing the networks. By incorporating principles of static and thermodynamic admissibility into the PINN framework, the proposed method significantly enhances the precision of material identification, especially in cases involving non-homogeneity.

The improvements are:

- Constraining the outputs of the neural network to ensure thermodynamically admissible constitutive behaviour.
- Enforcing static admissibility using the traditional Airy stress potential method.

Numerical experiments show that the enhanced PINN approach achieves accuracy levels comparable to traditional methods, such as CEGM, while offering significant reductions in computational time. In contrast, conventional PINNs without these physics-based constraints fail to reach similar levels of accuracy.

A key limitation of the current approach is the inherent smoothness of both the constitutive and stress fields. Compared to traditional constrained PINN methods, the Airy potential field has a weak instead of a strong discontinuity which is easier to approximate with NNs, but it is still overly smooth. While increasing the network size can improve approximation accuracy, it also leads to significantly higher computational costs. A more efficient alternative may be to use representations that inherently align with the continuity requirements of the constitutive, stress, or potential fields, potentially achieving greater accuracy at a lower computational expense.

Although this study is limited to small deformations and linear elasticity in two dimensions, the principle of enforcing static admissibility through stress potentials is broadly applicable. The requirement for a divergence-free Cauchy stress field is independent of the constitutive behaviour, making this approach viable for examples with large deformations and non-linear materials. While stress potentials can also be employed in three-dimensional problems, their formulation is significantly more complex. It remains an open question whether these advanced potential-based methods offer clear advantages in accuracy or efficiency over traditional PINNs.

The approach was further validated through real experiments, which confirmed accuracy and applicability of the method. It accurately identified the material properties of both homogeneous samples (nylon and TPU) and successfully determined the spatial distribution of material properties of the inhomogeneous specimens. The results closely matched the expected material properties reported in technical data sheets, while the identified boundaries of the inhomogeneous regions were in agreement with the actual geometry. These findings demonstrate the robustness of the method in real-world applications, where experimental data often includes noise and the fitted constitutive equation does not perfectly capture the material behaviour.

The main conclusion of this work is that physics considerations should play a central role in the design of PINNs. While NNs possess powerful learning capabilities, it is more effective to prescribe exact physical constraints where possible, reserving the use of NNs for solving parts that are challenging with traditional approaches. In this study, the constitutive behaviour and static admissibility requirements can be effectively described using traditional methods, while the NN is best suited for defining functional spaces and allows for the use of efficient and accelerated optimization algorithms.

CRediT authorship contribution statement

B. van der Heijden: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation. **X. Li:** Writing – review & editing, Formal analysis. **G. Lubineau:** Writing – review & editing, Conceptualization. **E. Florentin:** Writing – review & editing, Supervision, Methodology, Formal analysis, Conceptualization.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the authors used ChatGPT in order to improve the language and flow of the original draft. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the published article.

Funding

The authors acknowledge financial support received from King Abdullah University of Science and Technology (grant BAS/1/1315-01-01).

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

We want to thank Dr. Yi Wei for the kindness of sharing his Inverse-PINN code, and answering questions regarding his methodology that inspired this work.

Appendix. Normalization procedure

Characteristic values were defined to ensure that all normalized fields and parameters were approximately unity.[5] These characteristic properties are:

- **Characteristic length** l_0**:** is used to normalize the coordinate system,

$$\underline{\hat{x}} = \frac{\underline{x}}{l_0}. \tag{A.1}$$

- **Characteristic displacement** u_0**:** is used to normalize the displacement,

$$\underline{\hat{u}} = \frac{\underline{u}}{u_0}. \tag{A.2}$$

- **Characteristic force** F_0**:** is used to normalize the applied forces and tractions,

$$\underline{\hat{F}} = \frac{\underline{F}}{F_0} \tag{A.3}$$

In this study, the characteristic length was chosen as $l_0 = 10$ mm, $u_0 = \max(u_y) - \min(u_y)/5$, and $F_0 = F/5$. The latter two values varied between the three samples.

Normalization also affected the outputs and internal variables of the algorithm, so final results were scaled back using the characteristic properties. The following relationships were derived:

- Normalized strain scaling depends on length and displacement scaling:

$$\underline{\underline{\hat{\varepsilon}}} = \frac{\mathrm{d}\underline{\hat{u}}}{\mathrm{d}\underline{\hat{x}}} = \frac{\mathrm{d}\frac{\underline{u}}{u_0}}{\mathrm{d}\frac{\underline{x}}{l_0}} = \frac{l_0}{u_0}\underline{\underline{\varepsilon}}. \tag{A.4}$$

- Stress, is scaled with force and cross-sectional area:

$$\underline{\underline{\hat{\sigma}}} = \frac{\underline{\hat{F}}}{\hat{A}} = \frac{\frac{\underline{F}}{F_0}}{\frac{A}{l_0 t}} = \frac{l_0 t}{F_0}\underline{\underline{\sigma}}. \tag{A.5}$$

- Young's modulus, scaled with respect to stress and strain:

$$\hat{E} = \frac{\hat{\sigma}}{\hat{\varepsilon}} = \frac{u_0 t}{F_0}E. \tag{A.6}$$

---

[5] The hat notation $\hat{\square}$ indicates normalized quantities.

Data availability

No data was used for the research described in the article.

References

[1] P. Ladevéze, L. Chamoin, The constitutive relation error method: a general verification tool, in: Verifying Calculations - Forty Years on: An Overview of Classical Verification Techniques for FEM Simulations, Springer International Publishing, 2016, pp. 59–94, http://dx.doi.org/10.1007/978-3-319-20553-3_4.

[2] T. Kirchdoerfer, M. Ortiz, Data-driven computational mechanics, Comput. Methods Appl. Mech. Engrg. 304 (2016) 81–101, http://dx.doi.org/10.1016/j.cma.2016.02.001.

[3] N. Cottin, H.P. Felgenhauer, H.G. Natke, On the parameter identification of elastomechanical systems using input and output residuals, Ing.-Arch. 54 (5) (1984) 378–387, http://dx.doi.org/10.1007/BF00532820.

[4] K.T. Kavanagh, R.W. Clough, Finite element applications in the characterization of elastic solids, Int. J. Solids Struct. 7 (1) (1971) 11–23, http://dx.doi.org/10.1016/0020-7683(71)90015-1.

[5] D. Claire, F. Hild, S. Roux, A finite element formulation to identify damage fields: The equilibrium gap method, Internat. J. Numer. Methods Engrg. 61 (2) (2004) 189–208, http://dx.doi.org/10.1002/nme.1057.

[6] M. Grédiac, E. Toussaint, F. Pierron, L'identification des propriétés mécaniques de matériaux avec la méthode des champs virtuels, une alternative au recalage par éléments finis, Comptes Rendus Mécanique 330 (2) (2002) 107–112, http://dx.doi.org/10.1016/S1631-0721(02)01435-3.

[7] F. Pierron, M. Grédiac, The Virtual Fields Method: Extracting Constitutive Mechanical Parameters from Full-field Deformation Measurements, Springer New York, New York, NY, 2012, http://dx.doi.org/10.1007/978-1-4614-1824-5.

[8] G. Geymonat, F. Hild, S. Pagano, Identification of elastic parameters by displacement field measurement, Comptes Rendus Mécanique 330 (6) (2002) 403–408, http://dx.doi.org/10.1016/S1631-0721(02)01476-6.

[9] S. Avril, M. Bonnet, A.-S. Bretelle, M. Grédiac, F. Hild, P. Ienny, F. Latourte, D. Lemosse, S. Pagano, E. Pagnacco, F. Pierron, Overview of identification methods of mechanical parameters based on full-field measurements, Exp. Mech. 48 (4) (2008) 381–402, http://dx.doi.org/10.1007/s11340-008-9148-y.

[10] J.M.P. Martins, A. Andrade-Campos, S. Thuillier, Comparison of inverse identification strategies for constitutive mechanical models using full-field measurements, Int. J. Mech. Sci. 145 (2018) 330–345, http://dx.doi.org/10.1016/j.ijmecsci.2018.07.013.

[11] E. Florentin, G. Lubineau, Identification of the parameters of an elastic material model using the constitutive equation gap method, Comput. Mech. 46 (4) (2010) 521–531, http://dx.doi.org/10.1007/s00466-010-0496-y.

[12] S. Cuomo, V.S. Di Cola, F. Giampaolo, G. Rozza, M. Raissi, F. Piccialli, Scientific machine learning through physics–Informed neural networks: Where we are and What's next, J. Sci. Comput. 92 (3) (2022) 88, http://dx.doi.org/10.1007/s10915-022-01939-z.

[13] M. Raissi, P. Perdikaris, G. Karniadakis, Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, J. Comput. Phys. 378 (2019) 686–707, http://dx.doi.org/10.1016/j.jcp.2018.10.045.

[14] E. Zhang, M. Dao, G.E. Karniadakis, S. Suresh, Analyses of internal structures and defects in materials using physics-informed neural networks, Sci. Adv. 8 (7) (2022) http://dx.doi.org/10.1126/sciadv.abk0644.

[15] Y. Wei, Q. Serra, G. Lubineau, E. Florentin, Coupling physics-informed neural networks and constitutive relation error concept to solve a parameter identification problem, Comput. Struct. 283 (2023) 107054, http://dx.doi.org/10.1016/j.compstruc.2023.107054.

[16] A. Moussawi, G. Lubineau, E. Florentin, B. Blaysat, The constitutive compatibility method for identification of material parameters based on full-field measurements, Comput. Methods Appl. Mech. Engrg. 265 (2013) 1–14, http://dx.doi.org/10.1016/j.cma.2013.06.003.

[17] A. Moussawi, G. Lubineau, J. Xu, B. Pan, A 3D domain decomposition approach for the identification of spatially varying elastic material parameters, Internat. J. Numer. Methods Engrg. 102 (7) (2015) 1431–1448, http://dx.doi.org/10.1002/nme.4853.

[18] G.H. Lubineau, A. Moussawi, J. Xu, R. Gras, A domain decomposition approach for full-field measurements based identification of local elastic parameters, Int. J. Solids Struct. 55 (2015) 44–57, http://dx.doi.org/10.1016/j.ijsolstr.2014.11.009.

[19] P. Feissel, O. Allix, Modified constitutive relation error identification strategy for transient dynamics with corrupted data: The elastic case, Comput. Methods Appl. Mech. Engrg. 196 (13–16) (2007) 1968–1983, http://dx.doi.org/10.1016/j.cma.2006.10.005.

[20] S. Huang, P. Feissel, P. Villon, Modified constitutive relation error: An identification framework dealing with the reliability of information, Comput. Methods Appl. Mech. Engrg. 311 (2016) 1–17, http://dx.doi.org/10.1016/j.cma.2016.06.030.

[21] H.N. Nguyen, L. Chamoin, C. Ha Minh, mCRE-based parameter identification from full-field measurements: Consistent framework, integrated version, and extension to nonlinear material behaviors, Comput. Methods Appl. Mech. Engrg. 400 (2022) 115461, http://dx.doi.org/10.1016/j.cma.2022.115461.

[22] 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, Internat. J. Numer. Methods Engrg. 125 (8) (2024) e7439, http://dx.doi.org/10.1002/nme.7439.

[23] D.P. Kingma, J. Ba, Adam: A method for stochastic optimization, in: Y. Bengio, Y. LeCun (Eds.), 3rd International Conference on Learning Representations, San Diego, California, USA, 2015, pp. 1–15.

[24] Z. Aldirany, R. Cottereau, M. Laforest, S. Prudhomme, Multi-level neural networks for accurate solutions of boundary-value problems, Comput. Methods Appl. Mech. Engrg. 419 (2024) 116666, http://dx.doi.org/10.1016/j.cma.2023.116666.

[25] K. Hornik, M. Stinchcombe, H. White, Multilayer feedforward networks are universal approximators, Neural Netw. 2 (5) (1989) 359–366, http://dx.doi.org/10.1016/0893-6080(89)90020-8.

[26] A.K. Sarma, S. Roy, C. Annavarapu, P. Roy, S. Jagannathan, Interface PINNs (I-PINNs): A physics-informed neural networks framework for interface problems, Comput. Methods Appl. Mech. Engrg. 429 (2024) 117135, http://dx.doi.org/10.1016/j.cma.2024.117135.

[27] J.C. Platt, A.H. Barr, Constrained differential optimization, in: Proceedings of the 1987 International Conference on Neural Information Processing Systems, in: NIPS'87, MIT Press, Cambridge, MA, USA, 1987, pp. 612–621.

[28] W.S. Slaughter, Constitutive Equations, in: The Linearized Theory of Elasticity, Birkhäuser Boston, Boston, MA, 2002, pp. 193–220, http://dx.doi.org/10.1007/978-1-4612-0093-2.

[29] X. Glorot, Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in: Y.W. Teh, M. Titterington (Eds.), Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, in: Proceedings of Machine Learning Research, vol. 9, PMLR, Chia Laguna Resort, Sardinia, Italy, 2010-05-13/2010-05-15, pp. 249–256.

[30] A.F. Bower, Analytical techniques and solutions for linear elastic solids, in: Applied Mechanics of Solids, CRC Press, Boca Raton, 2009, pp. 248–380, http://dx.doi.org/10.1201/9781439802489.

[31] W.S. Slaughter, Forces and Stress, in: The Linearized Theory of Elasticity, Birkhäuser Boston, Boston, MA, 2002, pp. 157–192, http://dx.doi.org/10.1007/978-1-4612-0093-2.

[32] K.T. Chau, X.X. Wei, Stress Concentration Reduction at a Reinforced Hole Loaded by a Bonded Circular Inclusion, J. Appl. Mech. 68 (3) (2001) 405–411, http://dx.doi.org/10.1115/1.1357869.

[33] S.L. Omairey, P.D. Dunning, S. Sriramula, Development of an ABAQUS plugin tool for periodic RVE homogenisation, Eng. Comput. 35 (2) (2019) 567–577, http://dx.doi.org/10.1007/s00366-018-0616-4.

[34] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, A. Lerer, Automatic differentiation in PyTorch, in: NIPS Workshop on AutoDiff, Long Beach, California, USA, 2017, pp. 1–4.

[35] F. Latourte, A. Chrysochoos, S. Pagano, B. Wattrisse, Elastoplastic behavior identification for heterogeneous loadings and materials, Exp. Mech. 48 (4) (2008) 435–449, http://dx.doi.org/10.1007/s11340-007-9088-y.

[36] Ultimaker, TPU95A Technical Data Sheet, 2022, https://support.ultimaker.com/s/article/1667410781018.

[37] Ultimaker, Nylon Techincal Data Sheet, 2022, https://support.ultimaker.com/s/article/1667410781976.

[38] A. Yankin, Y. Alipov, A. Temirgali, G. Serik, S. Danenova, D. Talamona, A. Perveen, Optimization of Printing Parameters to Enhance Tensile Properties of ABS and Nylon Produced by Fused Filament Fabrication, Polymers 15 (14) (2023) 3043, http://dx.doi.org/10.3390/polym15143043.