

Blind multi-frequency identification of a localized complex modulus in an Euler–Bernoulli beam using variational physics-informed neural networks

M. Auger, R. Tittarelli, M. Ouisse

Université Marie et Louis Pasteur, SUPMICROTECH, CNRS, Institut FEMTO-ST
Besançon, F-25000, France
e-mail: maxime.auger@femto-st.fr

Abstract

This work addresses blind identification of a localized complex Young’s modulus in a clamped–free Euler–Bernoulli beam from partial complex frequency-response measurements. A multi-frequency variational physics-informed neural network jointly learns a frequency-conditioned displacement field and a frequency-independent complex modulus represented by smooth trainable plateaus. The beam equation is imposed on the observation window through a Petrov–Galerkin weak residual, while training and checkpoint selection remain oracle-free. Numerical tests at six frequencies from 500 to 3000 Hz recover both patch interfaces within the millimetre range and identify the patch loss modulus accurately. The low-loss background is recovered mainly through an effective mean level, while its left/right distribution remains less identifiable.

1 Introduction

Vibration damping is central to slender-structure design. Surface-bonded and constrained viscoelastic treatments increase dissipation with limited architectural change. Their characterization typically relies on modal or frequency-response measurements and mechanical models; the Oberst method, ASTM E756, and classical structural-damping references provide established routes [1].

For layered and constrained-layer beams, analytical models such as Ross–Kerwin–Ungar (RKU) describe equivalent stiffness and loss factors [2, 3, 4, 5]. Inverse use becomes harder when the dissipative region is unknown, measurements cover only part of the structure, and forcing or supports lie outside the observed window.

Here a hidden dissipative patch must be localized and quantified from complex displacement data on a subinterval, with excitation outside that window. The local weak equation is homogeneous, but the measured field still depends on unknown boundary traces, so the inverse problem is only partially identifiable. The added information consists of measured displacements, weak-form physics, multi-frequency consistency, and a low-dimensional modulus parameterization—not material oracle data.

Physics-informed neural networks combine data and governing equations [6] and have been used for beam parameter identification [7]. For Euler–Bernoulli beams, strong-form PINNs require high-order derivatives. Variational PINNs (VPINNs) instead impose a Petrov–Galerkin weak residual [8, 9], transferring derivatives to known test functions and making the formulation attractive for partial windows; accuracy still depends on test functions and quadrature [10, 11].

The first contribution is a blind multi-frequency VPINN in which a frequency-conditioned complex displacement network shares one frequency-independent complex Young’s modulus across all measured responses. This shared-modulus constraint regularizes the admissible profiles.

Second, the unknown modulus is represented by a small number of contiguous smooth plateaus with trainable amplitudes and transition locations. The two transitions around the patched region are therefore learned rather than prescribed.

Third, checkpoint selection is oracle-free: neither optimization nor selection uses true interfaces, FE reference moduli, or segmentation. Selection relies on hidden displacement validation, weak residuals on an independent shadow test bank, real/imaginary and local/global diagnostics, and recent parameter variability. For the oscillatory but improving trajectory studied here, the selector is interpreted as a diagnostic compromise criterion rather than evidence of unique stationary convergence.

Figure 1 summarizes the configuration: a clamped-free beam is excited upstream of a downstream observation window containing a hidden viscoelastic patch. The VPINN uses complex displacements and a weak Euler–Bernoulli residual on this window to identify a shared complex modulus. Section 2 defines the inverse problem, Section 3 the VPINN and selector, Section 4 the numerical results, and Section 5 the identifiability limits.

2 Problem setup

In the studied configuration, a prismatic Euler–Bernoulli beam of length 1 m is clamped at $x = 0$ and free at $x = 1m$. The coordinate x denotes the axial position along the beam. A harmonic point force of magnitude $F = 1N$ is applied at the point-load position

$$x_p = 0.25m, \quad (1)$$

which lies outside the observation window. The measured window is

$$\Omega_x = [x_{\min}, x_{\max}] = [0.5m, 0.995m]. \quad (2)$$

A viscoelastic patch is present in the FE reference model on the patched region

$$\Omega_{\text{patch}} = [0.7m, 0.8m]. \quad (3)$$

The word *patch* refers to this localized dissipative treatment, and *patched region* refers to the interval occupied by it. The *background* denotes the non-patched parts of the observation window,

$$\Omega_x \setminus \Omega_{\text{patch}} = [x_{\min}, 0.7m) \cup (0.8m, x_{\max}], \quad (4)$$

or, equivalently, the material contribution outside the viscoelastic treatment. The background modulus is the complex modulus associated with these non-patched regions.




Figure 1: Partial-window inverse problem. The point force lies outside the observation window; the hidden viscoelastic patch is used only for post-training assessment.

The complex Young's modulus is denoted

$$\tilde{E}(x) = E_r(x) + iE_i(x), \quad i = \sqrt{-1}, \quad i^2 = -1, \quad (5)$$

where $E_r(x)$ is the storage modulus and $E_i(x)$ is the loss modulus. The measured complex displacement at frequency f_ℓ is denoted

$$U_\ell(x) = U_{r,\ell}(x) + iU_{i,\ell}(x). \quad (6)$$

The displacement field is frequency-dependent, while the identified modulus $\tilde{E}(x)$ is assumed to be independent of frequency over the considered band. This assumption is central to the multi-frequency formulation: all measured frequency responses share the same unknown material profile.

The patch location and the material values are known only in the FE model used to generate reference data and in the post-run evaluation. They are not supplied to the inverse algorithm. In particular, the training procedure does not receive the true internal interfaces at $x = 0.7m$ and $x = 0.8m$, the patch width, the patch modulus, the background modulus, or a label indicating which points belong to the patch.

Since the excitation point x_p lies outside Ω_x , the force term is zero on the observation window. Nevertheless, the displacement field restricted to Ω_x depends on the upstream excitation and support through unknown boundary traces at the entrance of the window. These traces are not prescribed. Consequently, the local inverse problem cannot be made unique from the homogeneous differential equation alone. The proposed strategy therefore combines four ingredients: displacement data on Ω_x , a weak residual on Ω_x , multi-frequency consistency through a shared modulus, and a low-dimensional bubble representation of the material field.

3 Blind multi-frequency VPINN formulation

Following the physical setting introduced in Section 2, the inverse model is formulated on the observation window only. The unknown forcing and support effects outside this window are not prescribed; instead, their influence is absorbed through the frequency-conditioned displacement surrogate, while the material field is constrained through a shared weak-form residual.

3.1 Frequency-domain Euler–Bernoulli equation

For a harmonic excitation at frequency f_ℓ , the angular frequency is

$$\omega_\ell = 2\pi f_\ell. \quad (7)$$

The complex transverse displacement amplitude $U_\ell(x)$ satisfies

$$\frac{d^2}{dx^2} \left(\tilde{E}(x) I \frac{d^2 U_\ell}{dx^2} \right) - \rho S \omega_\ell^2 U_\ell(x) = -F \delta(x - x_p), \quad (8)$$

where S is the cross-sectional area, I is the second moment of area, ρ is the density, F is the point-force amplitude, x_p is the point-load position, and δ is the Dirac distribution. Primes or explicit derivatives with respect to x denote spatial derivatives, for example $U_\ell'' = d^2 U_\ell / dx^2$. On Ω_x , the right-hand side of equation (8) vanishes because $x_p \notin \Omega_x$.

The complex displacement and modulus are decomposed as

$$U_\ell(x) = U_{r,\ell}(x) + iU_{i,\ell}(x), \quad \tilde{E}(x) = E_r(x) + iE_i(x). \quad (9)$$

The bending-moment factor $\tilde{E}U_\ell''$ then reads

$$\tilde{E}U_\ell'' = (E_r U_{r,\ell}'' - E_i U_{i,\ell}'') + i(E_r U_{i,\ell}'' + E_i U_{r,\ell}''). \quad (10)$$

Thus the loss modulus E_i couples the real and imaginary equations: it contributes to the real part through $-E_i U_{i,\ell}''$ and to the imaginary part through $E_i U_{r,\ell}''$.

3.2 Weak residual on the observation window

The normalized coordinate $s \in [0,1]$ is defined by

$$x = x_{\min} + Js, \quad J = x_{\max} - x_{\min}. \quad (11)$$

Let $\phi_j(s)$ be test functions satisfying the endpoint conditions

$$\phi_j(0) = \phi_j(1) = 0, \quad \frac{d\phi_j}{ds}(0) = \frac{d\phi_j}{ds}(1) = 0. \quad (12)$$

These conditions remove the boundary terms generated by integration by parts. In practice, the test functions are constructed as

$$\phi_j(s) = b(s)g_j(s), \quad b(s) = s^2(1-s)^2, \quad (13)$$

where g_j is selected from a prescribed test-function family.

Multiplying the homogeneous form of equation (8) by ϕ_j over Ω_x and integrating twice by parts gives the Petrov–Galerkin weak residual

$$\mathcal{R}_{\ell j}(U_\ell, \tilde{E}) = \int_{\Omega_x} \tilde{E}(x) I U_\ell''(x) \phi_j''(x) dx - \int_{\Omega_x} \rho S \omega_\ell^2 U_\ell(x) \phi_j(x) dx. \quad (14)$$

Because of the conditions in equation (12), the unknown traces at the ends of the observation window do not enter explicitly in this weak residual. The integration by parts transfers part of the differentiation from the displacement surrogate to the known test functions, which is the main VPINN advantage for the Euler–Bernoulli operator.

The real and imaginary projected residuals are denoted $R_{k,\ell,j}$, with $k \in \{r, i\}$, frequency index ℓ , and test-function index j . They are

$$R_{r,\ell,j} = I \int_{\Omega_x} (E_r U_{r,\ell}'' - E_i U_{i,\ell}'') \phi_j'' dx - \rho S \omega_\ell^2 \int_{\Omega_x} U_{r,\ell} \phi_j dx \quad (15)$$

$$R_{i,\ell,j} = I \int_{\Omega_x} (E_r U_{i,\ell}'' + E_i U_{r,\ell}'') \phi_j'' dx - \rho S \omega_\ell^2 \int_{\Omega_x} U_{i,\ell} \phi_j dx. \quad (16)$$

In normalized coordinates, $\partial_x = J^{-1} \partial_s$ and $\partial_{xx} = J^{-2} \partial_{ss}$. Thus the bending contribution is scaled by I/J , whereas the inertial contribution is scaled by $\rho S \omega_\ell^2 J$ when the integrals are evaluated on $s \in [0,1]$.

The two terms entering each residual can have very different magnitudes depending on frequency and on the selected test function. For this purpose, each residual is divided by a scale built from its bending and inertia contributions. Define

$$A_{r,\ell,j} = I \int_{\Omega_x} (E_r U_{r,\ell}'' - E_i U_{i,\ell}'') \phi_j'' dx, B_{r,\ell,j} = \rho S \omega_\ell^2 \int_{\Omega_x} U_{r,\ell} \phi_j dx \quad (17)$$

$$A_{i,\ell,j} = I \int_{\Omega_x} (E_r U_{i,\ell}'' + E_i U_{r,\ell}'') \phi_j'' dx, B_{i,\ell,j} = \rho S \omega_\ell^2 \int_{\Omega_x} U_{i,\ell} \phi_j dx. \quad (18)$$

The normalized residuals are

$$\hat{R}_{k,\ell,j} = \frac{R_{k,\ell,j}}{(\langle A_{k,\ell,j}^2 \rangle_j^{1/2} + \langle B_{k,\ell,j}^2 \rangle_j^{1/2} + \varepsilon_R)}, \quad k \in \{r, i\}, \quad (19)$$

where $\varepsilon_R > 0$ is a small stabilizing constant. The notation $\langle \cdot \rangle_j$ denotes the average over the test functions used for the corresponding frequency and residual component. This normalization prevents a single frequency or test function from dominating the weak residual only because of scale. It is a numerical normalization and does not introduce information on the true modulus.

3.3 Frequency-conditioned displacement and shared modulus

Let

$$\mathcal{F} = \{f_1, \dots, f_{N_f}\} \quad (20)$$

be the set of measured frequencies. The frequency variable supplied to the displacement network is the normalized frequency

$$v_\ell = \frac{f_\ell - \bar{f}}{\sigma_f}, \quad (21)$$

where $\bar{f}$ and σ_f are the mean and standard deviation of the frequency set. The complex displacement surrogate is

$$U_\theta(s, v_\ell) = U_{r,\theta}(s, v_\ell) + iU_{i,\theta}(s, v_\ell), \quad (22)$$

where the two network outputs are the real and imaginary displacement components. The complex modulus surrogate is

$$\tilde{E}_\eta(s) = E_{r,\eta}(s) + iE_{i,\eta}(s). \quad (23)$$

It is not conditioned on v_ℓ . It is shared by all frequencies and therefore must be consistent with all measured responses.

The measured displacements are normalized frequency by frequency for the data loss. For each frequency,

$$U_\ell^{\max} = \max_m |U_\ell(x_m)|, \quad U_{r,\ell}^n(x_m) = \frac{U_{r,\ell}(x_m)}{U_\ell^{\max}}, \quad U_{i,\ell}^n(x_m) = \frac{U_{i,\ell}(x_m)}{U_\ell^{\max}}. \quad (24)$$

The superscript n denotes normalized displacement data. The data loss is evaluated in these normalized variables, whereas the weak residual in equations (15)–(16) is evaluated in physical units.

3.4 Bubble representation and smooth plateau functions

The material field is represented by a small number M of smooth plateaus. This global parameterization is called the *bubble representation*. Each elementary function in the representation is called a *smooth plateau function*. Let

$$0 \leq a_0 < a_1 < \dots < a_M \leq 1 \quad (25)$$

be the trainable edges in normalized coordinates. The smooth plateau function associated with segment m is

$$B_m(s) = \text{sigmoid}(\kappa(s - a_{m-1})) - \text{sigmoid}(\kappa(s - a_m)), \quad m = 1, \dots, M, \quad (26)$$

where $\kappa > 0$ controls the sharpness of the transitions. Thus each bubble corresponds to a smooth plateau between two learned transitions.

The storage and loss moduli are written as

$$E_{r,\eta}(s) = E_0 \sum_{m=1}^M A_{r,m} B_m(s), \quad E_{i,\eta}(s) = E_0 \sum_{m=1}^M A_{i,m} B_m(s) \quad (27)$$

where E_0 is a fixed scaling modulus and $A_{r,m}$ and $A_{i,m}$ are trainable dimensionless amplitudes. In the log-positive mode used for the numerical results,

$$A_{r,m} = 10^{\alpha_{r,m}}, \quad A_{i,m} = 10^{\alpha_{i,m}}. \quad (28)$$

This enforces positivity of both storage and loss moduli without prescribing their values.

The fixed modelling choices are the number of plateaus M , the functional form in equation (26), the modulus scale E_0 , and the transition sharpness κ . The learned quantities are the transition positions a_m , the plateau amplitudes $A_{r,m}$ and $A_{i,m}$, and the displacement-network parameters.

A neutral geometry regularization prevents degenerate segments but does not favour the true patch position. With $\Delta a_m = a_m - a_{m-1}$,

$$\mathcal{L}_{\text{geom}} = \lambda_g \left[\frac{1}{M} \sum_{m=1}^M \max(0, \Delta_{\min} - \Delta a_m)^2 + \frac{1}{M} \sum_{m=1}^M \max(0, \Delta a_m - \Delta_{\max})^2 + r_{\text{edge}} \right]. \quad (29)$$

Here λ_g is the regularization weight, $\Delta_{\min}$ and $\Delta_{\max}$ are minimum and maximum admissible normalized segment widths, and r_{edge} softly keeps the first and last edges close to the boundaries $s = 0$ and $s = 1$ of the measured window. This term uses only the measurement coordinates $x_{\min}$ and $x_{\max}$, not the true internal patch limits.

3.5 Training loss and adaptive real/imaginary weighting

For each frequency, the measured displacement points are split into a training set $\mathcal{D}_{\text{tr}_\ell}$ and a hidden validation set $\mathcal{D}_{\text{val}_\ell}$. The optimizer uses only $\mathcal{D}_{\text{tr}_\ell}$. The validation set $\mathcal{D}_{\text{val}_\ell}$ is hidden during optimization and is used only by the checkpoint selector.

The normalized data loss is

$$\mathcal{L}_{\text{data}}(\theta) = \frac{1}{N_f} \sum_{\ell=1}^{N_f} \frac{1}{|\mathcal{D}_{\text{tr}_\ell}|} \sum_{x_m \in \mathcal{D}_{\text{tr}_\ell}} \left[\left(U_{r,\theta}^n(s_m, v_\ell) - U_{r,\ell}^n(x_m) \right)^2 + \left(U_{i,\theta}^n(s_m, v_\ell) - U_{i,\ell}^n(x_m) \right)^2 \right]. \quad (30)$$

The weak residual losses are computed from the normalized projected residuals defined in equation (19):

$$\mathcal{L}_{\text{phys}}^r = \frac{1}{N_f N_t} \sum_{\ell=1}^{N_f} \sum_{j=1}^{N_t} \hat{R}_{r,\ell j}^2, \quad \mathcal{L}_{\text{phys}}^i = \frac{1}{N_f N_t} \sum_{\ell=1}^{N_f} \sum_{j=1}^{N_t} \hat{R}_{i,\ell j}^2. \quad (31)$$

The full objective is

$$\mathcal{J}(\theta, \eta) = w_d \mathcal{L}_{\text{data}} + w_r \alpha_r \mathcal{L}_{\text{phys}}^r + w_i \alpha_i \mathcal{L}_{\text{phys}}^i + \mathcal{L}_{\text{geom}}. \quad (32)$$

The weights w_d , w_r , and w_i set the nominal contributions of the data loss and of the real and imaginary weak residuals. The geometry term is controlled by λ_g in equation (29); equivalently, it may be interpreted as a separate weight w_g multiplying the geometry penalty.

The factors α_r and α_i are adaptive numerical weights. They are introduced to prevent one residual component from dominating the gradients only because of scale. Let p be the parameter set used for balancing, and let $\|\cdot\|_2$ denote the Euclidean norm over the gradients with respect to this parameter set. At a prescribed update frequency, i.e. every fixed number of optimizer iterations after an initial warm-up, the gradient magnitudes are evaluated as

$$g_r = \|\nabla_p \mathcal{L}_{\text{phys}}^r\|_2, \quad g_i = \|\nabla_p \mathcal{L}_{\text{phys}}^i\|_2, \quad g_* = (g_r g_i)^{1/2}. \quad (33)$$

The logarithmic targets are then

$$\log \alpha_r^* = \log(g_* + \varepsilon_g) - \log(g_r + \varepsilon_g), \quad \log \alpha_i^* = \log(g_* + \varepsilon_g) - \log(g_i + \varepsilon_g), \quad (34)$$

where $\varepsilon_g > 0$ stabilizes the logarithms. The targets are tracked by an exponential moving average and clipped to prescribed bounds. This adaptive weighting is a numerical rebalancing procedure. It uses no information on the true modulus or patch geometry.

3.6 Test functions for training and shadow diagnostics

The functions g_j in equation (13) can be chosen from several families. Local test functions, such as Gaussian radial basis functions (RBF), inverse multiquadrics (IMQ), and compactly supported Wendland functions, probe the residual over localized portions of the observation window. Global functions, such as Fourier and Legendre functions, probe coherence over the entire window. The two groups are not assumed to be universally better or worse; they provide complementary views of the weak residual.

Table 1 lists the families used in the numerical study. All functions are multiplied by the endpoint envelope $b(s) = s^2(1-s)^2$ so that the conditions in equation (12) hold.

Table 1: Test-function families used in the VPINN weak residual; each prototype is multiplied by the endpoint envelope.

Family	Prototype $g_j(s)$
Fourier	$\cos(k_j \pi s + \varphi_j), \sin(k_j \pi s + \varphi_j)$
Legendre	$P_{n_j}(2s - 1)$
Gaussian RBF	$\exp[-(s - c_j)^2 / (2\sigma_j^2)]$
Inverse multiquadric (IMQ)	$\left(1 + \epsilon_j^2(s - c_j)^2\right)^{-1/2}$
Wendland C^2	$(1 - \rho_j)_+^3(3\rho_j + 1)$

An independent shadow test bank, generated separately from the training bank, is used only after optimization. Shadow residuals are weak residuals evaluated at stored checkpoints on these unseen test functions. Local shadow residuals aggregate RBF, IMQ, and Wendland families; global residuals aggregate Fourier and Legendre families. They support the oracle-free checkpoint selector without updating the network.

3.7 Oracle-free checkpoint selector

The training trajectory is sampled at discrete checkpoints. A *checkpoint* k denotes a stored set of trainable parameters (θ_k, η_k) at a given cumulative iteration. The selector chooses one checkpoint after training without using the true modulus, the true interfaces, or the true segmentation.

The selector evaluates hidden displacement validation, shadow residuals on an independent test bank, real/imaginary and local/global residual components, and recent variability of interfaces and plateau amplitudes. Predefined oracle-free criteria identify a late admissible region, from which a representative checkpoint is selected rather than simply taking the final or minimum-loss iterate. Auxiliary candidates are retained to assess criterion sensitivity.

Here the selector is a diagnostic for an oscillatory but globally improving trajectory; it does not certify convergence or uniqueness. Its role is to avoid retaining a locally unfavourable late checkpoint and to compare physically meaningful candidates.

Hidden displacement validation. For checkpoint k , define

$$D_k = \frac{1}{N_f} \sum_{\ell=1}^{N_f} \frac{\|U_{\theta_k}^n(\mathcal{D}_{\text{val}_\ell}) - U_\ell^n(\mathcal{D}_{\text{val}_\ell})\|_2}{\|U_\ell^n(\mathcal{D}_{\text{val}_\ell})\|_2 + \varepsilon_D}, \quad (35)$$

where $\varepsilon_D > 0$ prevents division by zero. Real and imaginary components, D_k^r and D_k^i , are also stored. The validation data are displacement data only; they do not contain material labels.

Shadow weak residuals. Let q index a family in the shadow bank $\mathcal{B}_{\text{sh}}$. The selector computes

$$P_k^{(q,r)} = \frac{1}{N_f N_t^{(q)}} \sum_{\ell,j} \left(\hat{R}_{r,\ell j}^{(q)} \right)^2, \quad P_k^{(q,i)} = \frac{1}{N_f N_t^{(q)}} \sum_{\ell,j} \left(\hat{R}_{i,\ell j}^{(q)} \right)^2. \quad (36)$$

The local and global shadow residuals are then

$$P_k^{\text{loc}} = \text{median}_{q \in \{\text{RBF}, \text{IMQ}, \text{Wendland}\}} \left(P_k^{(q,r)} + P_k^{(q,i)} \right) \quad (37)$$

$$P_k^{\text{glob}} = \text{median}_{q \in \{\text{Fourier}, \text{Legendre}\}} \left(P_k^{(q,r)} + P_k^{(q,i)} \right). \quad (38)$$

A strict shadow score P_k^{sh} is formed from local, global, maximum-family, and real/imaginary-balance terms. Its role is to prevent a checkpoint with a small local residual but poor global consistency, or with a strong imbalance between real and imaginary residuals, from being selected.

Reduced variability of the learned modulus. The learned-modulus descriptor q_k contains only quantities produced by the inverse model:

$$q_k = [a_0, \dots, a_M, \log_{10} \bar{E}_{r,1}, \dots, \log_{10} \bar{E}_{r,M}, \log_{10} \bar{E}_{i,1}, \dots, \log_{10} \bar{E}_{i,M}]. \quad (39)$$

Here $\bar{E}_{r,m}$ and $\bar{E}_{i,m}$ are core medians on plateau m , computed away from the smooth transition layers. The recent variability over a window of W previously stored checkpoints is

$$S_k = \text{median}|q_k - \bar{q}_{k,W}|, \quad \bar{q}_{k,W} = \text{median}\{q_{k-1}, \dots, q_{k-W}\}. \quad (40)$$

A small S_k indicates that the learned interfaces and amplitudes vary slowly over the recent trajectory. This quantity is a relative diagnostic along the training path, not a proof that the inverse solution is unique or fully stationary.

Selector-defined admissible region. Because diagnostics may remain noisy after the main modulus features are recovered, a plateau hit is treated as an operational flag rather than proof of asymptotic convergence. It requires predefined low-slope, low-validation, low-shadow, learned-modulus-variability, and frequency-balance conditions.

Let $m_D(k)$ and $m_P(k)$ be rolling absolute slopes of $\log_{10} D_k$ and $\log_{10} P_k^{\text{sh}}$. Let B_k be the frequency-balance ratio, defined as the ratio between the worst validation error over frequencies and the mean validation error. A plateau hit satisfies

$$m_D(k) < \tau_D, \quad m_P(k) < \tau_P, \quad S_k < \tau_S, \quad B_k < B_{\max}. \quad (41)$$

Several consecutive plateau hits define a plateau-ready checkpoint and a late admissible region. The selected checkpoint is a representative compromise from this region.

Candidate scores. The selector computes several scores, all from oracle-free quantities. To compare metrics with different scales inside the training trajectory, each positive metric m_k is transformed into a robust logarithmic score

$$z_k(m) = \frac{\log_{10}(m_k + \varepsilon_z) - \text{median}_j \log_{10}(m_j + \varepsilon_z)}{1.4826 \text{median}_j |\log_{10}(m_j + \varepsilon_z) - \text{median}_p \log_{10}(m_p + \varepsilon_z)| + \varepsilon_z}. \quad (42)$$

For each metric, negative robust scores are better than the trajectory median and positive scores are worse; the median-absolute-deviation scaling only normalizes diagnostics along the run.

For a set of metrics $\mathcal{M}$ and non-negative weights β_m , define

$$C_k(\mathcal{M}, \beta) = \frac{\sum_{m \in \mathcal{M}} \beta_m z_k(m)}{\sum_{m \in \mathcal{M}} \beta_m}. \quad (43)$$

The balanced score is written in this form as

$$C_k^{\text{bal}} = C_k(\{D, D_{\max}, B_D, P^{\text{sh}}, P^{\text{loc}}, P^{\text{glob}}, B_P, S, G, m_D, m_P\}, \beta^{\text{bal}}). \quad (44)$$

The balanced score combines average and worst validation, real/imaginary and frequency balance, shadow residuals, learned-modulus variability, outer-plateau smoothness, and rolling slopes. Its weights are fixed before post-run assessment; the score is a selection compromise, not a claim of equal physical identifiability.

Two auxiliary scores use the same oracle-free metrics but emphasize imaginary terms (loss-sensitive) or real terms (storage-sensitive); neither uses reference modulus values.

The *early-stable score* is

$$C_k^{\text{early}} = C_k^{\text{bal}} + \beta_{\text{early}} \frac{k - k_{\min}}{k_{\max} - k_{\min} + \varepsilon_z}, \quad (45)$$

The early-stable score favours an earlier checkpoint within the candidate pool. A low-loss score combines the training objective with validation and shadow terms and is reported only as an auxiliary diagnostic. The primary checkpoint is the balanced candidate in the first late admissible region.

4 Numerical study

4.1 Numerical setup

In the studied numerical case, the beam is prismatic with length $L = 1m$, rectangular width $b = 0.02m$, and thickness $e = 0.002m$. The cross-sectional area and second moment of area are

$$S = be = 4e - 5m^2, \quad I = \frac{be^3}{12}. \quad (46)$$

The harmonic point force has magnitude $F = 1N$ and is applied at $x_p = 0.25m$. Six complex frequency-response functions are generated at

$$\mathcal{F} = \{500, 1000, 1500, 2000, 2500, 3000\} \text{ Hz}. \quad (47)$$

The complex displacement data are generated by the FE reference model and extracted on the observation window Ω_x . The viscoelastic treatment is represented in the FE model through an equivalent complex modulus following constrained-layer/RKU-type modelling [2, 3, 4]. The reference quantities introduced in this subsection are used for data generation and for the error assessment reported below; they are not part of the inverse inputs.

The words *steel* and *background* refer to the non-patched regions of the window, whereas *visco* and *patch* refer to the treated region Ω_{patch} . The FE reference loss moduli are

$$E_{i,FE}^{\text{patch}} = 1.025 \times 10^{11} \text{ Pa}, \quad E_{i,FE}^{\text{bg}} = 1.1275 \times 10^9 \text{ Pa}. \quad (48)$$

The same FE model supplies the storage-modulus references used in Tables 2 and 3 and in the dashed overlays of Figure 5.

For each frequency, 100 complex displacement samples are retained on Ω_x , giving 600 complex observations in total. The split is 80 training points and 20 hidden validation points per frequency. The displacement surrogate is a fully connected multilayer perceptron with three hidden layers of 64 neurons and hyperbolic tangent activations. Its inputs are s and the normalized frequency ν_ℓ ; its outputs are the normalized real and imaginary displacement components.

The modulus uses $M = 3$ smooth plateaus with $\kappa = 300$ and the log-positive parameterization in equation (28). The initial edges are uniform on $[0,1]$, which corresponds to approximately

$$x \simeq 0.50005m, 0.66503m, 0.83002m, 0.99500m. \quad (49)$$

This initialization does not use the true patch geometry.

The optimization is performed in two phases over a total of 150000 iterations. The first phase runs from iteration 0 to iteration 100000; it anchors the frequency-conditioned displacement field and gradually activates the weak residual. The second phase runs from iteration 100000 to iteration 150000; it refines the shared modulus and weak residuals with a lower learning rate and a larger weak-test space. In the reported implementation, Gauss–Legendre quadrature is used for the weak residual, with the training test family set to IMQ for the final run. The independent shadow bank contains Fourier, Legendre, Gaussian RBF, IMQ, and Wendland functions. The selector evaluates checkpoints every 500 cumulative iterations from iteration 10000 to the end of the run.

The plateau detector uses the following oracle-free thresholds in the reported run: $\tau_D = \tau_P = 2.5 \times 10^{-2}$ for the rolling log-slope conditions, $\tau_S = 8 \times 10^{-2}$ for the learned-modulus variability metric, and $B_{\text{max}} = 6$ for the frequency-balance ratio. A plateau-ready checkpoint requires four consecutive plateau hits. The quantile gates for validation and shadow metrics are 0.40 and 0.50, respectively. These thresholds define the selector admissibility criterion; they should not be interpreted as demonstrating complete stabilization of all diagnostics.

4.2 Blind multi-frequency identification results

Before presenting the numerical values, we define the reported error metrics. The interface errors are

$$x_1 - 0.7 \quad \text{and} \quad x_2 - 0.8, \quad (50)$$

reported in millimetres, where x_1 and x_2 are the two learned internal interfaces. The modulus errors are post-run relative errors,

$$e(E_k^{\text{patch}}) = 100 \frac{\bar{E}_k^{\text{patch}} - E_{k,\text{FE}}^{\text{patch}}}{E_{k,\text{FE}}^{\text{patch}}}, \quad e(E_k^{\text{bg}}) = 100 \frac{\bar{E}_k^{\text{bg}} - E_{k,\text{FE}}^{\text{bg}}}{E_{k,\text{FE}}^{\text{bg}}}, \quad k \in \{r, i\}. \quad (51)$$

Here $\bar{E}_k^{\text{patch}}$ and $\bar{E}_k^{\text{bg}}$ are core means or medians of the learned modulus in the corresponding regions, as used in the reported run. The displacement error ε_U is the global relative complex displacement error over all measured frequencies, reported in percent.

Table 2 compares the selected, final, and minimum-training-loss checkpoints; the primary checkpoint is the balanced candidate from the late admissible region.

Table 2: Blind multi-frequency identification results. The background is the union of the two non-patched regions of Ω_x .

Checkpoint	Iter.	$x_1 - 0.7$ (mm)	$x_2 - 0.8$ (mm)	$e(E_r^{\text{patch}})$ (%)	$e(E_r^{\text{bg}})$ (%)	$e(E_i^{\text{patch}})$ (%)	$e(E_i^{\text{bg}})$ (%)	ε_U (%)	Selection rule
Primary selector	146000	+0.455	-0.803	+1.012	+0.075	-0.326	-0.807	0.338	balanced admissible
Final	150000	+0.344	-0.493	+1.459	+0.309	+0.096	+0.408	0.320	end of training
Best loss	147990	+0.512	-0.637	+0.628	+0.152	+0.686	-0.878	0.321	minimum training loss

For the primary selected checkpoint, the identified loss moduli are

$$E_i^{\text{patch}} = 1.0217 \times 10^{11} \text{ Pa}, \quad E_i^{\text{bg}} = 1.1184 \times 10^9 \text{ Pa}. \quad (52)$$

The corresponding errors are -0.326% for the patch loss modulus and -0.807% for the mean background loss modulus. The two learned internal interfaces are $x_1 = 0.700455m$ and $x_2 = 0.799197m$. The global displacement error is 0.338% .

Figure 2 summarizes selector diagnostics. Real and imaginary validation errors become comparable and nearly flat late in training. Local shadow residuals drop after the phase transition but remain noisy, while global residuals stay larger. The diagnostics therefore indicate an admissible late region rather than complete weak-form stabilization.




Figure 2: Oracle-free selector diagnostics. Real/imaginary displacement validation and local/global shadow residuals are shown; markers indicate plateau hits, plateau-ready checkpoints, and the selected checkpoint.

Figure 3 shows the oracle-free candidate scores along training. Lower values are more favourable within each score, but the curves continue to fluctuate late in the run and do not share a common minimum. The selected checkpoint is therefore the balanced candidate in the late admissible region, where validation is low and the learned modulus is comparatively less variable.




Figure 3: Oracle-free candidate scores during training. Lower values are preferred within each score; the balanced selected checkpoint lies in the late admissible region.

Figure 4 shows the plateau diagnostics used by the selector. Rolling slopes and learned-modulus variability remain noisy, so plateau-hit and plateau-ready flags identify an admissible late region rather than a fully stationary inverse plateau.




Figure 4: Plateau diagnostics for the oracle-free selector. Markers identify plateau-hit, plateau-ready, and selected checkpoints; the noisy slopes indicate only partial stabilization.

Figure 5 shows the recovered modulus for the loss-focused auxiliary candidate. The patch interfaces and patch loss modulus are accurately recovered, whereas the two background plateaus retain a left/right ambiguity.




Figure 5: Recovered complex modulus for the loss-focused candidate. Interfaces and patch loss are accurately recovered, while the left/right background-loss split remains less constrained.

4.3 Candidate analysis and identifiability

Table 3 compares auxiliary candidates from the same late-training region to assess which identified quantities are robust to the selection criterion.

Table 3: Auxiliary oracle-free selector candidates. All candidates come from the same optimization trajectory.

Candidate	Iter.	$x_1 - 0.7$ (mm)	$x_2 - 0.8$ (mm)	$e(E_r^{\text{patch}})$ (%)	$e(E_r^{\text{bg}})$ (%)	$e(E_i^{\text{patch}})$ (%)	$e(E_i^{\text{bg}})$ (%)	ε_U (%)
Balanced / primary	146000	+0.455	-0.803	+1.012	+0.075	-0.326	-0.807	0.338
E_i -focused	134000	+0.242	-0.094	+1.041	+0.324	+0.479	-4.582	0.341
E_r -focused	140000	+0.250	+0.644	+1.069	+0.049	-0.977	+14.089	0.331
Low loss	145500	+0.466	-0.861	+1.017	+0.079	-0.238	-1.855	0.326

All late candidates identify the dissipative patch accurately. The internal interfaces and the patched loss modulus vary little from one candidate to another. The larger differences concern the low background loss modulus. The E_i -focused checkpoint is earlier and gives excellent interface localization, but a lower mean background loss than the balanced candidate. The E_r -focused checkpoint gives an accurate storage-modulus compromise but a larger relative error on the low background loss. The low-loss checkpoint remains close to the primary selected checkpoint.

The late candidates agree closely on the patch interfaces and patch loss modulus, whereas the two outer background plateaus can compensate each other. Their mean loss level may remain accurate while the left/right values differ. With partial-window data, weak background damping, and unknown entrance traces, the patch is therefore more identifiable than a unique background split.

5 Identifiability and limitations of the partial-window inverse problem

Using one shared modulus for six frequency responses improves conditioning: a profile that can be partly compensated by the displacement surrogate and unknown boundary traces at one frequency must remain consistent across all frequencies. In the present case, this chiefly improves identification of the patch interfaces and patch loss modulus.

Local and global shadow residuals provide complementary diagnostics. Local test functions are sensitive to localized modulus changes, whereas Fourier and Legendre functions probe weak-form consistency over the full observation window and therefore include the effect of unknown boundary traces. Their absolute levels need not be directly comparable. Because VPINN accuracy depends on test functions and quadrature [10, 11], the selector keeps local/global and real/imaginary diagnostics separate before combining them.

The real/imaginary split is likewise informative: imaginary quantities are directly sensitive to dissipation, although storage and loss moduli remain coupled through equation (10). The selected solution therefore balances displacement validation, weak residuals, learned-parameter variability, and real/imaginary consistency rather than fitting the imaginary response alone.

The diagnostics also delimit the selector's role. The trajectory improves globally but remains oscillatory: candidate scores drift, plateau-slope metrics stay noisy, and global shadow residuals remain above local ones. Because the final and minimum-loss checkpoints are also accurate, the selector is best viewed as an oracle-free safeguard against an unfavourable late iteration and as a tool for comparing plausible compromises.

Across these candidates, the patch interfaces and patch loss modulus are robust, while the detailed left/right background-loss distribution is not. This is consistent with the partial-window setting, where external excitation and support effects enter through unknown traces.

The discussion on background homogeneity must be treated cautiously. A background-homogeneity prior would be a physical regularization hypothesis stating that the two non-patched regions are made of the same material. It would not prescribe the numerical value of the background modulus. Such a prior is appropriate only if the experimental or manufacturing context supports the assumption that the two background regions are physically homogeneous. A less intrusive alternative is a post-training amplitude refit after interface stabilization. Once the learned transitions are fixed, the weak residual is approximately linear in the plateau amplitudes of E_r and E_i . A small regularized least-squares refit could then target the remaining amplitude ambiguity without using FE reference values.

Experimental validation remains necessary. Measurements will add sensor noise, excitation and geometry uncertainty, uncertain effective boundary traces, and model discrepancy. Future tests should therefore quantify robustness on measured frequency-response functions under controlled noise and displacement uncertainty.

6 Conclusion

This paper introduced a blind multi-frequency VPINN for identifying a localized complex modulus from partial complex displacement data. It combines a frequency-conditioned displacement surrogate, a shared frequency-independent modulus, a Petrov–Galerkin weak residual, and smooth trainable plateau functions.

The numerical study shows that six frequency responses from 500 Hz to 3000 Hz are sufficient to recover the two internal patch interfaces within the millimetre range and to identify the patch loss modulus with sub-percent accuracy. The primary oracle-free selector checkpoint at iteration 146000 yields a -0.326% error on the patched loss modulus and a -0.807% error on the mean background loss modulus. The final and best-loss checkpoints are also accurate. This indicates that, for the reported run, the optimization reaches a favourable late-training region; the remaining difficulty is the oscillatory character of the convergence, not a lack of global improvement.

The checkpoint selector is a diagnostic rather than proof of stabilization or uniqueness. It combines hidden displacement validation, independent shadow residuals, real/imaginary and local/global diagnostics, and recent variability of the learned parameters to select and compare representative late-training candidates.

The main limitation is background identifiability: the patch and its interfaces are strongly recovered, but the low background loss is constrained mainly through an effective mean level rather than a unique left/right split. This is consistent with the partial observation window and unknown entering traces.

Future work will target experimental validation, mixed local/global test spaces, and regularized post-training strategies for background identifiability, with robustness studies over patch locations, damping levels, noise, and neural-network initializations.

Acknowledgements

The authors acknowledge the support of EUR EIPHI, Contract No. ANR-17-EURE-0002. The authors also warmly thank the PINN team—Rafael Teloli, Patrice Le Moal, Valentin Calisti, and Emmanuel Ramasso—for fruitful discussions on physics-informed learning, inverse identification, and the methodological choices developed in this work.

References

- [1] A. D. Nashif, D. I. G. Jones, and J. P. Henderson, *Vibration Damping*. New York: John Wiley & Sons, 1985.
- [2] E. M. Kerwin, “Damping of flexural waves by a constrained viscoelastic layer,” *The Journal of the Acoustical Society of America*, vol. 31, no. 7, pp. 952–962, 1959.

- [3] D. Ross, E. E. Ungar, and E. M. Kerwin, "Damping of plate flexural vibrations by means of viscoelastic laminae," in *Structural Damping*. New York, NY: American Society of Mechanical Engineers, 1959, pp. 49–87.
- [4] D. J. Mead and S. Markus, "The forced vibration of a three-layer, damped sandwich beam with arbitrary boundary conditions," *Journal of Sound and Vibration*, vol. 10, no. 2, pp. 163–175, 1969.
- [5] D. J. Mead and S. Markus, "Loss factors and resonant frequencies of encastré damped sandwich beams," *Journal of Sound and Vibration*, vol. 12, no. 1, pp. 99–112, 1970.
- [6] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations," *Journal of Computational Physics*, vol. 378, pp. 686–707, 2019.
- [7] R. d. O. Teloli, R. Tittarelli, M. Bigot, L. Coelho, E. Ramasso, P. Le Moal, and M. Ouisse, "A physics-informed neural networks framework for model parameter identification of beam-like structures," *Mechanical Systems and Signal Processing*, vol. 224, p. 112189, 2025.
- [8] E. Kharazmi, Z. Zhang, and G. E. Karniadakis, "Variational physics-informed neural networks for solving partial differential equations," arXiv preprint, 2019. [Online]. Available: <https://arxiv.org/abs/1912.00873>
- [9] E. Kharazmi, Z. Zhang, and G. E. Karniadakis, "hp-vpinns: Variational physics-informed neural networks with domain decomposition," *Computer Methods in Applied Mechanics and Engineering*, vol. 374, p. 113547, 2021.
- [10] S. Berrone, C. Canuto, and M. Pintore, "Variational physics informed neural networks: The role of quadratures and test functions," *Journal of Scientific Computing*, vol. 92, no. 3, p. 100, 2022.
- [11] S. Berrone, C. Canuto, and M. Pintore, "Solving pdes by variational physics-informed neural networks: an a posteriori error analysis," arXiv preprint, 2022. [Online]. Available: <https://arxiv.org/abs/2205.00786>