Mesh-free equilibrium-regularised neural reconstruction of sparse displacement fields from digital volume correlation

Jan Mrógala[1,2,3,*], Irina Perfilieva[2]

[1]Department of Informatics and Computers, University of Ostrava, Ostrava, Czech Republic
[2]Institute for Research and Applications of Fuzzy Modeling, University of Ostrava, Ostrava, Czech Republic
[3]Czech Academy of Sciences, Institute of Geonics v. v. i., Ostrava, Czech Republic

Abstract

Digital Volume Correlation (DVC) measures the interior displacement field of a loaded specimen from a pair of micro-CT scans, but the measurement is sparse: where the image lacks texture, tracking fails and the displacement is unknown. We reconstruct the full field with a Physics-Informed Neural Network (PINN) that fits the sparse measurements while enforcing mechanical equilibrium (Navier–Cauchy, with stress by automatic differentiation) as a regulariser over the unobserved gaps. The network additionally receives the reference-image grey-level as an input, although on homogeneous materials an ablation shows this input to be neutral to slightly harmful so far. The result is a mesh-free, continuous, equilibrium-regularised reconstruction of a sparse DVC field, a neural counterpart to classical equilibrium-gap DVC. We select the architecture by a controlled study judged on error *in the gaps*. With gradient-norm loss balancing, the equilibrium constraint reduces mean gap-region error below a data-only baseline at every tested sparsity, with the largest margin where the data are sparsest, so it contributes most where observations are missing. On a real micro-CT rock-compression dataset the pipeline runs end-to-end in 3D, and on held-out nodes both trained fits rank above deterministic interpolants on average. The increment the equilibrium term adds over the data-only network there is, however, small and within run-to-run variability. By contrast, on the smooth, noise-free synthetic benchmark a thin-plate-spline interpolant remains the strongest method overall. We state explicitly under which conditions the physics term pays off, and the limitations of the present evidence.

Keywords

Physics-Informed Neural Networks, Digital Volume Correlation, Navier–Cauchy equation, Displacement measurement, Physics-based regularisation

CEUR-WS.org/Vol-4323/paper11.pdf

1. Introduction and problem

When a solid is loaded, every interior point moves; the vector field of those movements, the *displacement field*, is the primary quantity for understanding deformation and failure. DVC measures it by tracking local image texture between a reference and a deformed micro-CT scan [1]. The measurement is intrinsically *sparse and incomplete*: in homogeneous phases and pores the image carries no distinctive texture, correlation fails, and the displacement there is unknown, leaving gaps that must be filled before the field can be used.

A purely data-driven fit reproduces the measured displacements but is unconstrained in the gaps, where it may violate physics [2]. We fill the gaps with a PINN that adds *mechanical equilibrium* to the training objective as a regulariser, so the reconstruction is drawn toward mechanical admissibility everywhere, including unobserved regions. Spatial derivatives, and hence stress, come from automatic differentiation, so the trained network is a mesh-free, closed-form approximation of the displacement field. Our contribution occupies a gap in the literature (Section 2): *a mesh-free, equilibrium-regularised, continuous neural reconstruction of a sparse DVC displacement field, with the reference-image grey-level supplied as an auxiliary network input.* Section 2.1 motivates this formulation against directly solving the governing equations with an established solver; the experiments compare it against a matched

ITAT'26: Information Technologies – Applications and Theory, September 25–29, 2026, Biele Karpaty, Slovakia
*Corresponding author.
jan.mrogala@osu.cz (J. Mrógala); irina.perfilieva@osu.cz (I. Perfilieva)
© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

CEUR Workshop Proceedings ceur-ws.org ISSN 1613-0073
published 2026-11-09

data-only network *and* classical deterministic interpolants, and Section 6 states explicitly what the evidence does and does not support.

2. Related work and positioning

Two strands of prior work frame the contribution. **Classical mechanically-regularised / equilibrium-gap DVC** (Claire, Hild and Roux [3] and successors [4, 5, 6, 7]) augments correlation with the requirement that the recovered field satisfy equilibrium, acting as a mechanical filter. These methods are finite-element based: they solve a regularised linear system on a mesh and do not use the image itself as a model feature. **Neural DIC/DVC** reconstructs correlation with neural networks; PINN-DIC [8] is mesh-free but solves the correlation problem *from scratch via a photometric loss* and carries no mechanical equilibrium, while supervised convolutional DVC [9] predicts on a voxel grid from a labelled corpus and yields no closed-form derivatives or per-specimen physics.

Novelty relative to the two strands. To make the position explicit: relative to the equilibrium-gap lineage [3, 7] the method differs in representation, data, and inputs. (i) *Representation*: a coordinate network replaces the finite-element mesh, so the reconstruction is continuous and mesh-free, its strain and stress follow from automatic differentiation rather than element shape functions, and no mesh of the irregular, porous specimen is ever built. (ii) *Data*: the network consumes the scattered, already-correlated sparse field and treats completion as regularised regression, so it composes with any DVC front-end instead of modifying the correlation solver itself. (iii) *Inputs*: the reference-image grey-level enters as a model feature, which no equilibrium-gap formulation uses, preparing the ground for image-informed heterogeneous moduli (Section 6). Relative to PINN-DIC [8], which addresses the *correlation* problem itself and imposes no mechanical model, our network never sees the deformed image: it starts from DVC output and imposes equilibrium. The two are therefore complementary rather than competing – a PINN-DIC-style front-end could supply our input field.

2.1. Why optimise a network instead of solving the boundary-value problem?

Since the governing mechanical model is assumed known, it is fair to ask why the field should be reconstructed by training a neural network rather than by handing the problem to an established numerical solver. The answer is that the reconstruction task is *not a forward boundary-value problem.* A forward solve needs the specimen geometry meshed and complete boundary data on a closed boundary; here neither exists. The applied loads and platen contact tractions are not measured, and displacement data exist only at the scattered interior points where correlation converged – the gaps routinely extend to the specimen surface, so there is no closed boundary on which Dirichlet or Neumann data could be prescribed. Moreover, the measurements carry correlation noise: prescribing them as exact constraints of a forward solve would propagate that noise through the domain, which is precisely the failure mode for which the exact interpolants are penalised in Section 5. The task is therefore one of *data assimilation*: find a field that approximately satisfies equilibrium and matches noisy scattered observations in a least-squares sense. Classical mechanical regularisation poses it exactly this way – the equilibrium-gap method solves a regularised least-squares system on a finite-element mesh [3, 4] – and our formulation keeps that variational structure while replacing the finite-element space with a neural one. The meaningful comparison is thus not "PINN versus forward FEM" but "neural least squares versus finite-element least squares".

What the neural parameterisation buys is flexibility. It requires no mesh of a porous, irregular specimen, ingests scattered 3D points natively, and yields a continuous field with analytic derivatives, so strains are evaluated at arbitrary points without post-hoc numerical differentiation of noisy grids. Auxiliary inputs (here the grey-level) and constitutive extensions (spatially varying moduli, nonlinear laws) are added without re-deriving a discretisation, and the same optimisation framework extends to parameter identification, where material properties are trained jointly with the field [10, 11, 12]. What it costs

is compute and guarantees: minutes of gradient-based optimisation (Section 4) against seconds for a sparse linear solve, and no discretisation-error theory. The intended application is accordingly not to replace finite-element mechanics but to complete, denoise, and differentiate DVC measurements in settings where meshing is impractical – and, prospectively, to support image-informed identification of heterogeneous properties, for which the image input is already in place.

3. Method

3.1. Physics and loss

Reconstruction problem. The input is a reference micro-CT image, a regular grid (a matrix in 2D, a voxel array in 3D) of grey-levels $g_{\text{ref}}(\mathbf{X})$, together with the *sparse* displacement vectors $\mathbf{u}_i^{\text{meas}}$ that DVC recovers only at the textured points $\mathbf{X}_i$, a subset of the grid; at the remaining, untextured grid points the displacement is unknown. The desired output is the *complete* displacement field: a vector $\mathbf{u}(\mathbf{X})$ at every grid point, equal to the measurements where they exist and filling the gaps where they do not. We represent this field by a network $\mathbf{u}_\theta$, a continuous function, evaluated *pointwise*: at a grid location $\mathbf{X}$ it takes that location's coordinate concatenated with its grey-level $g_{\text{ref}}(\mathbf{X})$ and returns the displacement vector *at that same location.* The output vector at $\mathbf{X}$ is thus the reconstructed displacement of the material point imaged at pixel/voxel $\mathbf{X}$; evaluating $\mathbf{u}_\theta$ over the whole grid maps the input grey-level image to a displacement field of matching shape (two or three scalar components per grid point). Because $\mathbf{u}_\theta$ is a continuous function with smooth derivatives, the reconstruction is mesh-free and defined between grid points as well, which is what lets the equilibrium residual below be imposed at arbitrary collocation points, including inside the gaps.

We treat a homogeneous, isotropic, linear-elastic material. For displacement $\mathbf{u}(\mathbf{X})$ the small strain is $\boldsymbol{\varepsilon} = \frac{1}{2}(\nabla\mathbf{u} + \nabla\mathbf{u}^\top)$ and the stress is $\boldsymbol{\sigma} = \lambda\,\mathrm{tr}(\boldsymbol{\varepsilon})\mathbf{I} + 2\mu\boldsymbol{\varepsilon}$, with Lamé parameters from a fixed (E, ν). Mechanical equilibrium is the Navier–Cauchy equation $\nabla\cdot\boldsymbol{\sigma} + \mathbf{b} = \mathbf{0}$; for the equilibrium-satisfying ground truth we use $\mathbf{b} = \mathbf{0}$. The training loss combines a data term on the textured points and an equilibrium-residual term on collocation points over the *whole* domain (including the gaps):

$$\mathscr{L} = \frac{w_d}{N_d}\sum_i \|\mathbf{u}_\theta(\mathbf{X}_i) - \mathbf{u}_i^{\text{meas}}\|^2 + \frac{w_p}{N_c}\sum_j \|\nabla\cdot\boldsymbol{\sigma}(\mathbf{u}_\theta)(\mathbf{X}_j)\|^2. \tag{1}$$

The residual takes *second* derivatives of the output, so the network's activations must be at least C^2. Following the reference image, the network input is the coordinate $\mathbf{X}$ concatenated with the reference grey-level $g_{\text{ref}}(\mathbf{X})$; the output is the displacement vector (displacement-only). We non-dimensionalise coordinates and displacements to $O(1)$; because $\nabla\cdot\boldsymbol{\sigma} = \mathbf{0}$ is homogeneous in E for $\mathbf{b} = \mathbf{0}$, this leaves the (scale-invariant) relative-L^2 metrics unchanged while making the data/physics loss balance meaningful.

Homogeneous Navier–Cauchy with constant moduli is blind to microstructure, yet we feed the image as a feature; on homogeneous synthetic data the grey-level input may therefore be neutral or slightly harmful, which we quantify in an ablation below that motivates the heterogeneous-material extension.

3.2. Architecture selection

We fit one coordinate-MLP per specimen, a continuous map from coordinates and grey-level to displacement with analytic derivatives by autodiff, chosen over CNN/U-Net DVC (needs a labelled corpus; no closed-form derivatives) and neural operators (data-hungry). Three requirements drive the design and map onto known evidence: (1) the loss needs *accurate second derivatives*, favouring smooth or periodic (SIREN) representations [13]; (2) gap-filling is a *frequency* problem where spectral bias [14] over-smooths the gaps, which random Fourier features mitigate at a tuned scale σ [15, 16]; (3) competing loss terms plus the image branch cause gradient imbalance, addressed by loss-weight balancing

and an improved backbone [17]; sizing follows the wide-layer result [18, 19] in the PINN setting of Haghighat et al. [10].

We therefore do not assume an architecture but *select* it by a controlled study, varying one axis at a time with everything else fixed (same case, mask, collocation set, optimiser budget; three seeds, two for the σ/ω_0 pre-tune): Axis A activation/encoding {tanh, tanh+Fourier (σ tuned), SIREN (ω_0 tuned)}; Axis B backbone {plain, improved}; Axis C output {displacement-only, mixed $(\mathbf{u}, \boldsymbol{\sigma})$}; Axis D depth/width. Configurations are ranked by, in order: (i) gap-region relative-L^2 of $\mathbf{u}$; (ii) strain/second-derivative accuracy in the gaps; (iii) cross-seed stability; (iv) cost. On the smooth cantilever the study selects a plain TANH MLP (Table 1): high-frequency encodings hurt on this smooth field, where aggressive Fourier scales and large SIREN ω_0 fail to fit even the observed data and produce diverging strains. The refinements that matter are the training rather than the activation: an unbalanced physics weight over-weights equilibrium and degrades reconstruction, whereas gradient-norm annealing removes the hand-tuning, and the improved backbone [17] lowers gap error by an order of magnitude at fixed weight (gap relative-L^2 $0.0280 \to 0.0024$), both as the gradient-pathology evidence predicts. We use the TANH MLP with gradient-norm balancing.

4. Experimental setup

Synthetic ground truth. The primary case is the Timoshenko–Goodier end-loaded cantilever (plane stress), a closed-form displacement/stress field satisfying $\nabla \cdot \boldsymbol{\sigma} = \mathbf{0}$ with zero body force; we verify numerically that our autodiff stress matches the analytic Timoshenko stress and that the residual vanishes to machine precision. A procedural speckle field provides g_{ref}, and the data/gap mask is derived from its local texture, mimicking DVC failure in untextured zones.

Implementation, runtime, and reproducibility. All models are implemented in PyTorch and trained in float32 (the physics verification above runs in float64); the equilibrium residual takes second derivatives of the network output by automatic differentiation. Unless a study varies it, every run uses a 5-layer, 128-unit MLP trained with Adam (learning rate 10^{-3}) for 20 000 iterations on a 120×40 grid with $\nu = 0.22$ (only ν enters the non-dimensional residual); collocation points are all material grid points, including the gaps. Gradient-norm balancing follows Wang et al. [17]: every 100 iterations the data weight is pulled toward the ratio of the maximum gradient magnitude of the weighted physics term to the mean gradient magnitude of the data term, smoothed by an exponential moving average (factor 0.9), with the physics term acting as the anchor. Every configuration runs with three seeds (two for the σ/ω_0 pre-tune); training is deterministic, and a checkpointed restart reproduces the uninterrupted trajectory exactly, so all quoted numbers are reproducible from the stored run configurations. On one Graphics Compute Die of an AMD Instinct MI250X (LUMI), a 20 000-iteration data-only run takes ~47 s and an equilibrium-regularised run ~8 min – the second-order residual dominates the cost. Real-data runs use a 20 000-node subsample and 15 000 iterations; the reconstruction of Fig. 3 (25 000-node subsample giving 15 030 data and 26 290 collocation points, 4 000 iterations) trains in ~85 min on 16 CPU threads. The implementation, run configurations, and aggregated run outputs are available from the corresponding author on request. The gap metric is the relative L^2 error $\|\mathbf{u}_\theta - \mathbf{u}\|_2 / \|\mathbf{u}\|_2$ evaluated only at the withheld points – the model's held-out test error (its accuracy where no training data exist).

Validation by masking. As real DVC has no ground truth in the gaps, we use analytical fields with truth known everywhere, withhold the data in masked regions, and score against the withheld truth. The headline metric is the gap-region relative-L^2 of $\mathbf{u}$ against a data-only baseline ($w_p = 0$); we also ablate the grey-level input, sweep the weight ratio and balancing, and sweep the mask fraction. The protocol transfers to real data via held-out well-correlated SPAM points [20].

Table 1
Axis-A architecture study on the cantilever, all at the same reference configuration (30% withheld, grey-level on, w_p=1, static weighting): gap-region relative-L^2 of $\mathbf{u}$, mean±std over three seeds (lower better), each encoding at its best tuned frequency. The selected architecture is shown in bold.

architecture	tuned param	gap rel-L^2
tanh	–	**0.0280±0.0052**
FF	σ=0.5	0.0820±0.0303
SIREN	ω_0=5	0.1574±0.0225

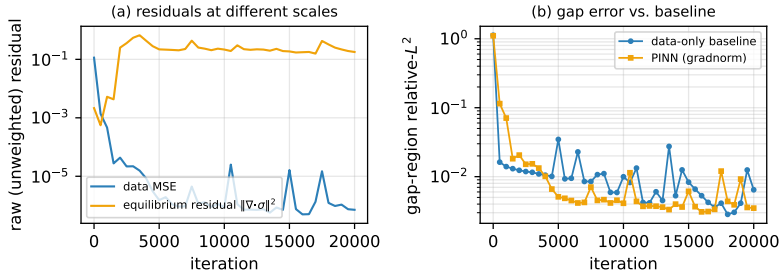


Figure 1: Training convergence (cantilever, 30% withheld, one seed). (a) Raw (unweighted) residuals. (b) Gap-region relative-L^2 during training: the gradient-norm-balanced PINN converges below the identical data-only baseline.

5. Results

The results below aggregate 72 synthetic configurations (196 runs in total; three seeds for every configuration quoted below) and 30 runs on the real dataset.

Architecture. Table 1 ranks Axis A: the plain TANH MLP wins on the smooth field, Fourier/SIREN encodings degrading as their frequency grows. The training choices matter more than the encoding (Section 3): gradient-norm balancing and the improved backbone turn an over-weighted, harmful physics term into a beneficial one.

Training behaviour. Figure 1 tracks a representative run: the (squared) data loss falls to $\sim 10^{-6}$ once the observations are fit, whereas the second-derivative equilibrium residual floors near 10^{-1} – an offset reflecting operator order, not violated equilibrium – while the balanced PINN's gap-region error falls below the identical data-only baseline and stays there.

Equilibrium regularisation vs. sparsity (headline). With gradient-norm balancing, the PINN's mean gap-region error is below the identical data-only baseline at *every* withholding fraction (by 17–60%), with the largest margin at the sparsest settings (Fig. 2): at 80% of the field withheld the PINN reaches gap relative-L^2 0.0038 versus 0.0096 for the baseline, a 60% reduction, while the baseline both worsens and scatters across seeds. (Both models are the plain TANH MLP; gap errors sit well below the reference cell of Table 1 because gradient-norm balancing replaces the over-weighted static w_p=1 used there.) Each seed poses the identical withheld field to both models, so the runs compare in seed-matched pairs: the PINN wins 12 of 15 (one-sided sign test p=0.018), including every pair at the two sparsest fractions; margins at mild sparsity sit within overlapping error bars (a consistent trend), while the sparsest setting is clearly separated – there the PINN's cross-seed std is also ~8× smaller than the baseline's. This is graceful degradation: the equilibrium constraint contributes most where observations are missing. Figure 2a also scores two deterministic interpolants under the same protocol: cubic interpolation loses to the PINN from 30% withheld upward, whereas a thin-plate-spline interpolant wins outright here – by more than an order of magnitude at mild withholding, narrowing

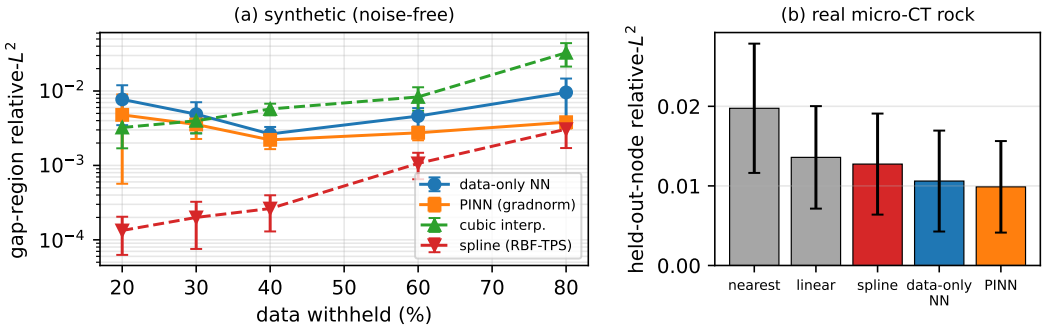


Figure 2: Comparison against deterministic interpolation (mean±std; dashed = interpolants). (a) Synthetic cantilever: the PINN degrades gracefully where the data-only network does not, but on this smooth, noise-free field a thin-plate-spline interpolant is the strongest method of all. (b) Real micro-CT rock, held-out reliable nodes: with measurement noise the ranking reverses and both trained fits rank above every interpolant on average.

to ~20% over the PINN at 80% withheld – since its smoothness prior matches this smooth, noise-free analytic field; on real, noisy measurements that ranking reverses on average (Fig. 2b). The synthetic study should therefore be read as isolating the effect of the equilibrium term against an otherwise identical network, not as a claim of superiority over classical interpolation on smooth data (Section 6).

Ablations. On these homogeneous synthetic fields the grey-level input is neutral to slightly harmful: a constitutive law with constant moduli is blind to microstructure, so the network cannot exploit an image feature yet, which motivates the heterogeneous-moduli extension in Section 6. The mixed $(\mathbf{u}, \boldsymbol{\sigma})$ output helps at large physics weight by halving the differentiation order of the residual.

Real micro-CT rock. On a SPAM-correlated compression dataset of one rock specimen (~61% of nodes converged, ~39% failed), the pipeline runs end-to-end in 3D. For this specimen the failed nodes lie predominantly in the background rather than the interior (correlation succeeded throughout the material), so we assess reconstruction by the masking protocol (Fig. 3): a coherent block of well-correlated data at the largest load step is withheld and reconstructed. The figure displays the fluctuation about the dominant near-affine trend; its structure – two counter-rotating in-plane circulation cells – is restored inside the block (0.062-voxel vector RMSE). On this smooth, well-correlated specimen a block of any placement or thickness we tried stays within reach of plain interpolation (linear: 0.019 voxels), so the figure demonstrates end-to-end reconstruction of real, structured data; the advantage appears where measurements are scattered and noisy. Scored on randomly held-out reliable nodes (three seeds × three load steps), the PINN reaches 0.0099±0.0058 versus 0.0106±0.0063 for the data-only baseline. This is a slight improvement that we deliberately do not overstate: averaging the three load steps within each seed gives three seed-matched pairs, and the PINN wins all three, but a one-sided sign test on three pairs cannot reach significance (p=0.125), so the increment is within the reported variability – unsurprisingly, since such held-out nodes are well surrounded by data. Deterministic interpolation of the same held-out nodes ranks behind both trained fits on average (linear 0.0136±0.0064, thin-plate spline 0.0127±0.0063, Fig. 2b): an exact interpolant reproduces the correlation noise of the neighbouring nodes, which the regularised fits average out. We do not claim statistical separation for this ordering either: at three seeds none of the pairwise orderings among the leading methods is resolved, and what the data support is the average ranking and its mechanism. The one unqualified real-data finding is that an *unbalanced* physics weight harms reconstruction here too, confirming that loss balancing is essential in practice; whether the equilibrium term itself buys a real-data accuracy margin remains open at this sample size.

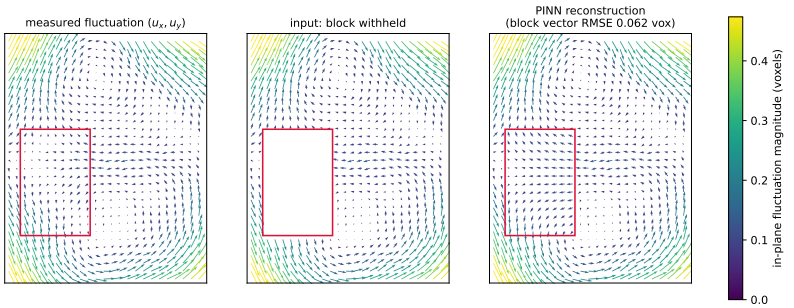


Figure 3: Real micro-CT rock, largest load step, one z-slice (masking protocol). Arrows: in-plane *fluctuation* (measured displacement minus its best-fit affine trend), coloured by magnitude – two counter-rotating circulation cells. The red block is withheld from training (centre) and reconstructed (right) at 0.062-voxel vector RMSE (linear interpolation: 0.019 voxels on this thin block).

6. Discussion and limitations

What is claimed, and what is not. On controlled synthetic data with known truth, the properly balanced equilibrium term reduces gap-region error below an otherwise identical data-only fit at every tested sparsity, and most where data are sparsest; this is the claim the study design isolates, and the paired statistics support it. It is *not* a claim of superiority over classical interpolation on smooth data: on the noise-free cantilever a thin-plate spline is the strongest method at every withholding fraction – by over an order of magnitude at mild sparsity (Fig. 2a) – because a generic smoothness prior is exactly the right prior for a smooth analytic field. The trained fits pull ahead, on average, only on the real, noisy measurements (Fig. 2b), where exact interpolation reproduces correlation noise. Likewise, on the real specimen the increment of the equilibrium term over the data-only network is small and within run-to-run variability; the defensible real-data statement is that the regularised *fits* rank above the interpolants on average, not that any pairwise ordering there is statistically resolved at three seeds, and not that the physics term has a demonstrated real-data margin.

Constitutive assumptions. The physics prior is homogeneous, isotropic, small-strain linear elasticity with zero body force. For rock this is a deliberate simplification: the material is heterogeneous at the imaged scale, and compaction or micro-cracking under load violates linear elasticity locally. We therefore use equilibrium strictly as a *regulariser* whose weight is balanced against the data – never as a hard constraint – so that where measurement and model disagree, the data term dominates. This mis-specification is also the most plausible reason the grey-level input does not yet pay off (a constant-moduli residual is blind to microstructure) and may contribute to the small real-data increment. Image-informed spatially varying moduli [12] are the direct remedy, and the reason the image input is kept in the architecture despite an ablation that is so far neutral to slightly harmful.

Breadth of the real-data evidence. All real-data results come from a single rock specimen, one instrument, and one correlation configuration (SPAM, 30-voxel node spacing) over three load steps. This specimen is, moreover, favourable to simple interpolation in contiguous-block tests (Fig. 3): roughly 94–98% of the raw displacement is explained by a global affine trend, and correlation succeeded throughout the material interior, so the held-out-node protocol necessarily tests reconstruction between well-surrounded points. We therefore claim no generality across materials, load regimes, or DVC settings. Specimens whose interior correlation genuinely fails – where no ground truth exists for scoring, and where the physics prior must bridge real gaps – are the target scenario; evaluating on the further specimens of the same experimental campaign is the immediate next step.

Cost, and when to prefer which method. The equilibrium-regularised fit costs minutes of optimisation where a spline costs seconds and a classical finite-element regularised solve is a sparse linear

system. The evidence here supports a plain recommendation: on dense, low-noise, smooth fields, use a spline; where a good mesh and boundary model exist, mature FE-based mechanical regularisation [7] is the established tool; the neural formulation earns its cost when the data are scattered, noisy, and three-dimensional, when meshing is impractical, when analytic strains of the reconstruction are needed, or when image-informed constitutive extensions and parameter identification are the goal.

7. Conclusion and future work

We presented a mesh-free, equilibrium-regularised neural reconstruction of sparse DVC displacement fields with the reference image as an auxiliary input, selected its architecture by an evidence-based controlled study, and evaluated it against a matched data-only network and classical deterministic interpolants on synthetic fields with known truth and on a real micro-CT compression experiment. The evidence supports three statements. First, with gradient-norm balancing the equilibrium term reduces gap-region error below the identical data-only fit at every tested sparsity (12/15 seed-paired wins), with the largest and most stable margin where the data are sparsest: the physics contributes most where observations are missing. Second, the decisive design choices are on the training side – loss balancing and the improved backbone – rather than exotic encodings, which on smooth fields hurt. Third, on real, noisy measurements both trained fits rank above every interpolant tested on average, while the increment of the equilibrium term over the data-only fit is small and within variability; on smooth, noise-free synthetic data a thin-plate spline remains the best method overall. The method therefore earns its place in noisy, sparse, mesh-hostile settings; it is not a universal replacement for interpolation (Section 6). Future work targets image-informed spatially varying moduli, which would turn the currently inert grey-level input into a constitutive signal; evaluation across further specimens; a systematic study of how accuracy depends on the size of a contiguous masked window (the spatial extent the physics prior must bridge unaided, not only the withheld fraction); and the extension toward fracture/phase-field modelling and parameter identification [21, 11].

Acknowledgments

Computational resources were provided by the LUMI supercomputer. The micro-CT sandstone data were obtained with support from the long-term conceptual development of the research organisation RVO: 68145535 (Institute of Geonics of the Czech Academy of Sciences). This work was supported by the Student Grant Competition (SGS) of the University of Ostrava.

Declaration on Generative AI

During the preparation of this work, the authors used Claude (Anthropic) in order to: Drafting content, Grammar and spelling check. After using these tool(s)/service(s), the authors reviewed and edited the content as needed and take full responsibility for the publication's content.

References

[1] A. Buljac, C. Jailin, A. Mendoza, J. Neggers, T. Taillandier-Thomas, A. Bouterf, B. Smaniotto, F. Hild, S. Roux, Digital volume correlation: Review of progress and challenges, Experimental Mechanics 58 (2018) 661–708.

[2] M. Raissi, P. Perdikaris, 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 378 (2019) 686–707.

[3] D. Claire, F. Hild, S. Roux, A finite element formulation to identify damage fields: the equilibrium gap method, International Journal for Numerical Methods in Engineering 61 (2004) 189–208.

[4] J. Réthoré, S. Roux, F. Hild, Extended and integrated digital image correlation: the equilibrium gap method as a mechanical filter, European Journal of Computational Mechanics 18 (2009) 285–306.
[5] H. Leclerc, J.-N. Périé, F. Hild, S. Roux, Voxel-scale digital volume correlation, Experimental Mechanics 51 (2011) 479–490.
[6] Z. Tomičević, F. Hild, S. Roux, Mechanics-aided digital image correlation, Journal of Strain Analysis for Engineering Design 48 (2013) 330–343.
[7] A. Mendoza, J. Neggers, F. Hild, S. Roux, Complete mechanical regularization applied to digital image and volume correlation, Computer Methods in Applied Mechanics and Engineering 355 (2019) 27–43.
[8] B. Li, S. Zhou, Q. Ma, S. Ma, PINN-DIC: A physics-informed neural network approach for digital image correlation, Experimental Mechanics (2024). ArXiv:2409.00956.
[9] X. Duan, J. Huang, Deep learning-based digital volume correlation, Extreme Mechanics Letters 53 (2022) 101710.
[10] E. Haghighat, M. Raissi, A. Moure, H. Gomez, R. Juanes, A physics-informed deep learning framework for inversion and surrogate modeling in solid mechanics, Computer Methods in Applied Mechanics and Engineering 379 (2021) 113741.
[11] D. Anton, H. Wessels, Physics-informed neural networks for material model calibration from full-field displacement data, 2022. ArXiv:2212.07723.
[12] S. Srikitrungruang, et al., IE-PINN: Identification of heterogeneous elastic properties from noisy displacement data, Advanced Science (2025). ArXiv:2506.14036.
[13] V. Sitzmann, J. Martel, A. Bergman, D. Lindell, G. Wetzstein, Implicit neural representations with periodic activation functions, in: Advances in Neural Information Processing Systems (NeurIPS), volume 33, 2020, pp. 7462–7473.
[14] N. Rahaman, A. Baratin, D. Arpit, F. Draxler, M. Lin, F. Hamprecht, Y. Bengio, A. Courville, On the spectral bias of neural networks, in: International Conference on Machine Learning (ICML), 2019, pp. 5301–5310.
[15] M. Tancik, P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. Barron, R. Ng, Fourier features let networks learn high frequency functions in low dimensional domains, in: Advances in Neural Information Processing Systems (NeurIPS), volume 33, 2020, pp. 7537–7547.
[16] S. Wang, H. Wang, P. Perdikaris, On the eigenvector bias of Fourier feature networks: From regression to solving multi-scale PDEs with physics-informed neural networks, Computer Methods in Applied Mechanics and Engineering 384 (2021) 113938.
[17] S. Wang, Y. Teng, P. Perdikaris, Understanding and mitigating gradient flow pathologies in physics-informed neural networks, SIAM Journal on Scientific Computing 43 (2021) A3055–A3081.
[18] A. Pinkus, Approximation theory of the mlp model in neural networks, Acta Numerica 8 (1999) 143–195.
[19] L. Lu, X. Meng, Z. Mao, G. E. Karniadakis, DeepXDE: A deep learning library for solving differential equations, SIAM Review 63 (2021) 208–228.
[20] O. Stamati, E. Andò, E. Roubin, R. Cailletaud, et al., spam: Software for practical analysis of materials, Journal of Open Source Software 5 (2020) 2286.
[21] S. Goswami, C. Anitescu, S. Chakraborty, T. Rabczuk, Transfer learning enhanced physics informed neural network for phase-field modeling of fracture, Theoretical and Applied Fracture Mechanics 106 (2020) 102447.