Inverse Physics-Informed Neural Network for CSFM Stress Field Reconstruction in Concrete D-Regions

Sandesh Lamsal * Rubi Bhandari

AtkinsRéalis USA Inc., 800 Waterford Way, Miami, FL 33126, USA
Department of Civil and Architectural Engineering, University of Miami, Coral Gables, FL 33146, USA
Department of Civil and Environmental Engineering, Florida International University, Miami, FL 33174, USA

Abstract

The Compatible Stress Field Method (CSFM) delivers the complete stress field of a structural-concrete discontinuity region (D-region) by combining a continuous stress field with kinematic compatibility and realistic, softening material laws. For an existing or monitored structure the internal stress state is the quantity of interest for assessment, but a finite-element analysis cannot supply it from measurements alone: it requires the loads and the full boundary data as input. This study formulates the recovery of the field as an inverse problem and solves it with a physics-informed neural network. Strain measured at a set of points is assimilated together with the CSFM physics, the cracked rotating compression field, kinematic compatibility and the traction boundary conditions, by a mixed formulation in which one network carries the strain field and a second the displacement field; the stress follows from a differentiable embedding of the cracked-membrane constitutive law. On a deep-beam D-region the internal stress field is reconstructed to within about 7% of an independent continuum CSFM solution from a dense, low-noise strain field, and on a cantilever-bracket corbel, a qualitatively distinct D-region with a single inclined strut and a top tie, to about 12%. Applied to the geometry of an experimentally tested wall pier from the validation suite of the CSFM, the reconstruction reaches 1.5%, though the simultaneously recovered axial and horizontal forces carry larger errors of 7.6% and 10.7% respectively. The error grows with measurement noise, to about 19% at 3% gauge noise on the deep beam, as the constitutive amplification quantified below predicts. This study identifies and quantifies the limiting mechanism: the cracked constitutive law amplifies strain error by a factor of about 3.4, because the crack direction is governed by the shear strain, so that measurement uncertainty propagates magnified into the stress. The applied load, never supplied to the network, is recovered to within 1.2% on the deep beam and 8.3% on the corbel by integrating the reconstructed vertical stress across an interior horizontal cut, where the gauge data densely constrains the field; the naive boundary integration on the load patch itself, where the field is least constrained, gives errors above 25% and is reported only as a diagnostic. A forward physics-only network is shown to be unsuitable, and the strong-form equilibrium residual is found to be detrimental as a training objective.

Keywords: inverse problem, physics-informed neural network, Compatible Stress Field Method, discontinuity regions, stress-field reconstruction, structural concrete.

1 Introduction

Reinforced-concrete structures contain regions in which the Bernoulli hypothesis of plane sections does not hold: the vicinity of concentrated loads and supports, frame corners, corbels, deep beams, openings, pile caps and hammerhead piers. In these *discontinuity regions* (D-regions) the strain distribution is markedly nonlinear, and design and assessment must follow the actual flow of forces. Strut and tie modelling, placed on a consistent footing for design

*Corresponding author: sandeshlamsal@miami.edu

practice by Schlaich et al. [21], idealises that flow as a pin-jointed truss and applies the lower-bound theorem of plasticity to deliver a safe design. The Compatible Stress Field Method (CSFM), developed by Kaufmann et al. [9] and now embedded in commercial design software, generalises this idea: it combines a continuous stress field with kinematic compatibility and finite, realistic material laws, a parabola-rectangle law for concrete in compression with the compression-softening of cracked concrete and a bilinear law for reinforcement. The CSFM thereby yields not a few member forces but the complete, continuous stress and strain field of a D-region, together with its deformation and failure load.

The companion study in this series [12] built a fast surrogate that predicts the failure load of a D-region on a fixed strut and tie graph. Graph-based neural surrogates of this kind have proved effective for real-time structural prediction more broadly, including the form-finding of anticlastic tensile membranes [11]. The present study turns to a different question, and one that the surrogate and the solver alike leave unanswered: given an existing or monitored D-region, what is its *internal* stress state? This is the quantity an engineer needs in order to assess a structure already in service, to interpret a load test, or to track a member instrumented for structural health monitoring, increasingly viable on real structures thanks to distributed-strain sensing technologies such as those reviewed by Glisic and Inaudi [5]. A finite-element stress-field analysis cannot answer it from measurements: the analysis is a *forward* operator that maps prescribed loads and boundary data to the field, and in an assessment setting the loads are precisely what is uncertain or unknown. What is available instead is the opposite, partial and indirect information: the structure stands in equilibrium under whatever it carries, and a limited number of sensors record its response.

Recovering the field from such information is an *inverse* problem, and it is here, rather than as a forward surrogate, that physics-informed learning has a genuine role. A physics-informed neural network (PINN) can assimilate sparse measurements together with the governing physics in a single variational problem, a setting in which it has no finite-element counterpart. A purely forward PINN, one trained only against the governing equations to reproduce the solver, is by contrast not competitive for this problem: the strong-form equilibrium residual it minimises admits a diffuse, near-equilibrated field far from the true one. The inverse problem does not share this weakness, because the measurements anchor the solution.

This study presents an inverse PINN for the CSFM. Strain measured at a set of points is assimilated with the CSFM physics to reconstruct the full internal stress field of a D-region. To make the reconstruction well conditioned the network is *mixed*: one network represents the strain field directly and a second the displacement field, the two coupled by kinematic compatibility. The cracked-membrane constitutive law of the CSFM, including the compression-softening of cracked concrete, is embedded as a differentiable map, so that the reconstructed stress is a CSFM stress field by construction. The mixed strain–displacement architecture, together with the finding that the strong-form equilibrium residual is detrimental as a training objective and that compatibility, the constitutive law and the traction conditions alone suffice, defines an inverse PINN for the CSFM from which the applied load is recovered to within a few per cent by integrating the reconstructed vertical stress across an interior horizontal cut. Beyond the reconstruction itself, the work isolates the mechanism that bounds the achievable accuracy: the cracked constitutive law amplifies strain error into stress error by a factor of about 3.4, an error budget mapped here against measurement density and noise and applicable to any strain-based CSFM assessment, not only to the network presented below.

This study is the continuum-field stage of a planned series on machine-learning methods for the CSFM; it builds on the notation and the reference solver of the companion discrete study [12].

2 Related work

Stress fields and the CSFM. The lower-bound, plasticity-based design of D-regions by trusses and stress fields was placed on a consistent footing by Schlaich et al. [21] and Marti [14], and generalised to continuous fields by Muttoni et al. [15]. The treatment of cracked concrete on which the CSFM depends originates in the modified compression-field theory of Vecchio and Collins [23], which established that the compressive strength of concrete falls as transverse cracking increases, and in the cracked membrane model of Kaufmann and Marti [8]. The CSFM of Kaufmann et al. [9] assembles these ingredients into a code-oriented method that delivers the stress field, the deformation and the failure load within one framework; its continuum form is the physics this study reconstructs.

Stress-field reconstruction from measured strain: the virtual fields method. Reconstructing the internal stress field from measured strain has a substantial history outside machine learning, principally through the virtual fields method of Pierron and Grédiac [17]. The virtual fields method exploits the principle of virtual work: for any kinematically admissible virtual displacement field $\mathbf{u}^*$ with associated virtual strain $\boldsymbol{\varepsilon}^*$, the experimental strain field $\boldsymbol{\varepsilon}^{\text{exp}}$ and the unknown stress field $\boldsymbol{\sigma}$ must satisfy $\int_\Omega \boldsymbol{\sigma} : \boldsymbol{\varepsilon}^* \, dV = \int_{\partial\Omega_t} \mathbf{b} \cdot \mathbf{u}^* \, dS$, which becomes a linear (or in nonlinear settings, nonlinear) system once $\boldsymbol{\sigma}$ is parameterised through a constitutive law. The method is the workhorse of inverse identification in experimental mechanics; it is mature, theoretically well founded, and routinely applied to full-field strain measurements from digital image correlation. The present study differs in three respects that are material for the CSFM application. First, the virtual fields method typically identifies a small number of constitutive parameters or a piecewise-low-order stress field; it does not by default deliver the full continuous σ_2 field that an inverse PINN parameterises with a network. Second, it requires the user to construct virtual fields that are sensitive to the quantity of interest and decoupled from the unknowns, a step that becomes delicate for a cracked, softening, direction-dependent constitutive law such as the CSFM's; the network formulation sidesteps this construction entirely. Third, the virtual fields method is designed for full-field measurements (DIC); it accommodates sparse gauge data only with auxiliary interpolation, while the network ingests scattered gauge points directly. The two methods are complementary and the inverse PINN should not be read as a replacement: the virtual fields method remains the right tool for parametric identification on full-field DIC data; the network is the right tool for full continuum-field reconstruction from sparse gauges under a cracked, nonlinear constitutive law.

Physics-informed neural networks and inverse problems. Physics-informed neural networks, introduced by Raissi et al. [18] and surveyed by Karniadakis et al. [7], embed the residual of a governing differential equation in the loss of a neural network. From the outset the framework was posed for two distinct tasks: the forward solution of a differential equation, and the *inverse* problem of inferring fields and parameters from sparse data constrained by the same physics. The inverse task is the one in which a PINN has no direct classical counterpart, because it fuses measurements and physics in a single variational problem; Haghighat et al. [6] demonstrate the approach on inverse identification problems in linear-elastic solid mechanics, recovering stress fields and material parameters from sparse strain or displacement data, and it is in this setting that the present study sits. A complementary strand, the deep energy method of Samaniego et al. [20], replaces the strong-form residual by the system's potential energy. Training such networks is not straightforward: Wang et al. [24] document gradient-flow pathologies that stall convergence when competing loss terms are unbalanced, which has motivated adaptive loss-weighting schemes such as that of Bischof and Kraus [4].

Physics-informed learning for structural concrete. PINNs have begun to be applied to structural members. Bastek and Kochmann [2] develop PINNs for shell structures, and, most directly related to this study, Balmer et al. [1] apply PINNs to the nonlinear analysis of reinforced-concrete beams. Their work is a one-dimensional Euler-Bernoulli beam with a multi-linear moment-curvature law; they introduce a mixed-PINN formulation, with separate networks for the deflection and the bending moment coupled by a loss term. Mixed formulations in solid-mechanics PINNs trace back further to the deep energy methods of Nguyen-Thanh et al. [16], where the displacement and deformation gradient are treated as separately parameterised fields coupled by compatibility. This study extends physics-informed learning for structural concrete in two respects: from the one-dimensional forward beam problem to the two-dimensional *inverse* reconstruction of a continuum stress field, and from a multi-linear law to the CSFM cracked compression field including compression softening. The mixed formulation of Balmer et al. [1] and the SiLU activation [19] are adopted here.

3 The continuum Compatible Stress Field Method

This section states the continuum CSFM in the form the inverse network must satisfy. A D-region is treated as a two-dimensional plane-stress domain Ω of constant out-of-plane thickness, loaded in its plane.

Kinematics. Let $\mathbf{u}(\mathbf{x}) = (u_x, u_y)$ be the in-plane displacement field. The small-strain tensor is its symmetric gradient, with components

$$\varepsilon_x = \frac{\partial u_x}{\partial x}, \quad \varepsilon_y = \frac{\partial u_y}{\partial y}, \quad \gamma_{xy} = \frac{\partial u_x}{\partial y} + \frac{\partial u_y}{\partial x}. \tag{1}$$

Cracked rotating compression field. The CSFM models the concrete as a cracked, rotating compression field: the principal stress and strain directions coincide, the concrete tensile strength is neglected, and reinforcement is smeared into the element with ratios ρ_x, ρ_y following the cracked membrane model [8]. From the strain $(\varepsilon_x, \varepsilon_y, \gamma_{xy})$ the principal strains $\varepsilon_1 \geqslant \varepsilon_2$ and the inclination θ of the principal frame follow in closed form,

$$\theta = \tfrac{1}{2}\,\mathrm{atan2}(\gamma_{xy},\, \varepsilon_x - \varepsilon_y)\,, \tag{2}$$

so that the crack direction is governed by the shear strain. Concrete in compression follows the parabola-rectangle law; the effective compressive strength of cracked concrete is reduced by the compression-softening factor

$$k_{c2}(\varepsilon_1) = \min\left(1,\ \frac{1}{0.8 + 140\,\varepsilon_1}\right), \tag{3}$$

in the form adopted by the CSFM of Kaufmann et al. [9] after the cracked-membrane model of Kaufmann and Marti [8], with the coefficient 140 matching their implementation; the original Vecchio–Collins MCFT [23] uses 170, and the difference reflects the distinct calibration of the cracked-membrane treatment of tension stiffening. Reinforcement follows an idealised bilinear law. The resulting map from strain to the global stress $\boldsymbol{\sigma} = (\sigma_x, \sigma_y, \tau_{xy})$, here denoted $\boldsymbol{\sigma} = \mathcal{C}(\varepsilon; \rho_x, \rho_y)$, is the constitutive heart of the method.

Equilibrium and boundary conditions. The stress field must satisfy plane-stress equilibrium in the interior,

$$\nabla \cdot \boldsymbol{\sigma} + \mathbf{b} = \mathbf{0} \quad \text{in } \Omega, \tag{4}$$

with $\mathbf{b}$ the body force (here negligible), the traction condition $\boldsymbol{\sigma} \cdot \mathbf{n} = \bar{\mathbf{t}}$ on the loaded and free parts of the boundary, and the prescribed displacement at the supports. A forward CSFM analysis prescribes the loads and solves Equations (1)–(4) for the field; the inverse problem of this study prescribes, instead of the loads, a set of strain measurements.

4 Inverse-PINN methodology

4.1 The inverse problem

A D-region is instrumented at M points $\{\mathbf{x}_m\}$ at which the in-plane strain is recorded,

$$\tilde{\varepsilon}_m = \boldsymbol{\varepsilon}(\mathbf{x}_m) + \boldsymbol{\eta}_m, \qquad m = 1, \ldots, M, \tag{5}$$

with $\boldsymbol{\eta}_m$ the measurement noise. The applied load is not known. The support locations and the traction-free parts of the boundary are known, as is the reinforcement layout and the material. The task is to reconstruct the full stress field $\boldsymbol{\sigma}(\mathbf{x})$ over Ω. Because the loads are absent, this is not a forward boundary-value problem; it is the assimilation of the measurements (5) with the physics of Section 3.

4.2 A mixed strain–displacement network

A single displacement network is not adequate. Its strain field, obtained as the symmetric spatial gradient of the network output, inherits a magnified fitting error, so that a displacement field matched to the data to one per cent can carry a strain field in error by tens of per cent. This study therefore adopts a mixed formulation, after Balmer et al. [1]. Two networks are trained jointly:

$$\boldsymbol{\varepsilon}(\mathbf{x}) \approx \mathcal{E}_\vartheta(\mathbf{x}), \qquad \mathbf{u}(\mathbf{x}) \approx \mathcal{U}_\phi(\mathbf{x}), \tag{6}$$

each a multilayer perceptron with the SiLU activation. The strain network $\mathcal{E}_\vartheta$ carries the strain field directly, as a network output rather than as a derivative, so that the measurements anchor it cleanly. The displacement network $\mathcal{U}_\phi$ supplies kinematic compatibility: the two are tied by requiring the strain network to equal the symmetric gradient of the displacement network. The stress is obtained from the strain network through the embedded constitutive map,

$$\boldsymbol{\sigma}(\mathbf{x}) = \mathcal{C}(\mathcal{E}_\vartheta(\mathbf{x}); \rho_x, \rho_y), \tag{7}$$

so the reconstruction is, by construction, a CSFM stress field. Figure 1 summarises the architecture.

4.3 Differentiable constitutive embedding

The cracked-membrane map $\mathcal{C}$ of Section 3 is ported to a differentiable, automatic-differentiation-traceable form. It was verified against the continuum CSFM solver: applied to the solver's own element strains it reproduces the solver's element stresses to within 0.6–2.7% (root-mean-square, relative), confirming a faithful port. The compression-softening factor (3) is smooth; the remaining slope discontinuities of the law (compression onset, steel yield) are mild.

4.4 Loss

The networks are trained by minimising

$$\mathcal{L} = w_d\, \mathcal{L}_{\text{data}} + w_c\, \mathcal{L}_{\text{comp}} + w_t\, \mathcal{L}_t + w_u\, \mathcal{L}_u. \tag{8}$$

The *data* term is the mismatch between the strain network and the measurements (5), normalised component by component (the vertical strain ε_y is several times smaller than the others, and a common normalisation would let the optimiser ignore it):

$$\mathcal{L}_{\text{data}} = \frac{1}{3} \sum_{c \in \{x,y,xy\}} \frac{\frac{1}{M} \sum_m \left(\mathcal{E}_{\vartheta,c}(\mathbf{x}_m) - \tilde{\varepsilon}_{m,c} \right)^2}{\langle \tilde{\varepsilon}_c^2 \rangle}. \tag{9}$$

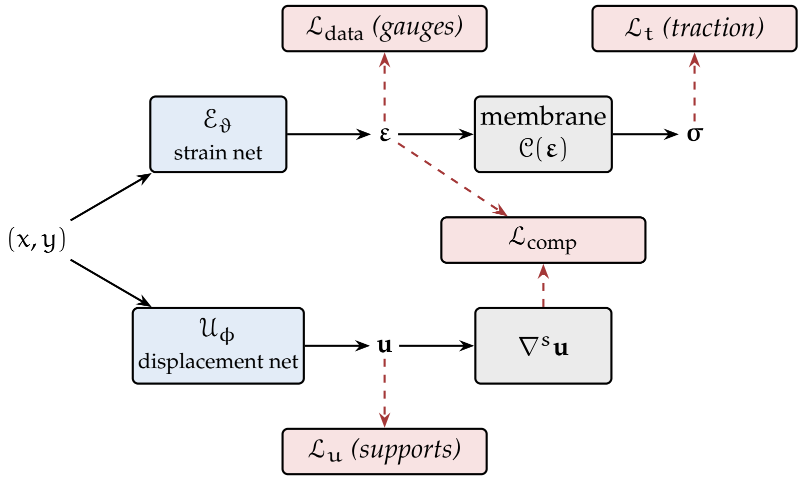

Figure 1: Architecture of the inverse PINN. The coordinate (x, y) feeds two networks: the strain network $\mathcal{E}_\vartheta$ outputs the strain field directly, and the displacement network $\mathcal{U}_\phi$ outputs the displacement. Stress is obtained by passing the strain network's output through the embedded CSFM cracked-membrane constitutive map $\mathcal{C}$. Solid arrows are field operations; dashed arrows are loss terms. The data loss $\mathcal{L}_{\text{data}}$ anchors the strain on the gauges; the compatibility loss $\mathcal{L}_{\text{comp}}$ ties the two networks ($\varepsilon \approx \nabla^s \mathbf{u}$); $\mathcal{L}_t$ enforces the traction boundary conditions and $\mathcal{L}_u$ the support conditions. The strong-form interior equilibrium residual is intentionally absent: it was found detrimental (Section 5).

The *compatibility* term enforces the mixed coupling, $\mathcal{L}_{\text{comp}} = \langle \|\mathcal{E}_\vartheta - \nabla^s \mathcal{U}_\phi\|^2 \rangle$, evaluated on interior collocation points; $\mathcal{L}_t$ collects the traction-free residual on the unloaded boundary and the condition that the load-patch traction is purely vertical; and $\mathcal{L}_u$ enforces the known support displacements. Equilibrium (4) is *not* included as a loss term. The strong-form equilibrium residual was found to be detrimental: enforced as a soft penalty it drives the reconstruction toward a spurious near-equilibrated field, exactly the pathology that defeats the forward PINN, and it degrades an otherwise accurate reconstruction (Section 5). Compatibility, the constitutive map and the traction conditions, with the data, are sufficient.

4.5 Training

Both networks are six-hidden-layer multilayer perceptrons with a width of 128 neurons per layer and SiLU activations. The loss weights in Equation (8) are set to $w_d = 10$, $w_c = 2$, $w_t = 1$ (applied with a half weight on the load-patch shear term), and $w_u = 5$; the weights are fixed across all archetypes and all sweeps. Adam runs for 3 000 iterations on data and compatibility, then 5 000 iterations with the traction and support terms active, at a base learning rate of 2×10^{-3}; 2 000 interior collocation points are resampled each step.

The networks are trained with the Adam optimiser [10]. When the measurements are noise-free the data term is driven down and an L-BFGS polish [13] follows. When the measurements carry noise, training is stopped once the data term reaches its noise floor, the value it attains when the network equals the true field and only the measurement noise remains; minimising past that floor over-fits the noise.

Computational cost. Training is performed on a workstation-class CPU (12 threads, PyTorch float32, no GPU acceleration) and takes 2.5–3 hours of wall-clock time per archetype, dominated by the 8 000 Adam iterations; the subsequent L-BFGS polish adds approximately 15 minutes. The bottleneck in our implementation is the serial Python overhead of the differentiable constitutive map (Section 3) rather than the network forward pass, so a GPU port offers limited speed-up without first vectorising the constitutive evaluation; porting the constitutive map to a vectorised

GPU kernel is straightforward and is left to the released code. Once trained, the network reconstructs the full continuous stress field at any density of query points in under a second per evaluation, and the interior-cut load recovery is essentially free. The cost profile is therefore that of an offline calibration (per geometry, per instrumented structure) followed by real-time online inference, well suited to the strain-based monitoring use case sketched in Section 1 (*cf.* Glisic and Inaudi 5): a network trained once on a structure's nominal continuum-CSFM solution can turn each new strain-measurement update into a fresh stress-field estimate at interactive rates, without re-solving the inverse problem from scratch.

4.6 Reference solution

The reconstruction is validated against the continuum CSFM solver. That solver discretises the D-region with constant-strain triangles, whose element stresses carry heavy element-scale discretisation noise: a 50 mm and a 25 mm mesh of the same problem differ by 49% element-wise in the minor principal stress. A constant-strain-triangle field is therefore not, as computed, a usable point-wise reference for a smooth reconstruction. The reference used here is a fine (12.5 mm) analysis with Gaussian stress recovery, a smoothing step in the spirit of finite-element stress-recovery practice; the recovered field is mesh-converged to within about 3–5% in the minor principal stress. Synthetic measurements are drawn from this reference field, with additive Gaussian noise, and the reconstruction is validated against it.

5 Results

The method is demonstrated on a simply supported deep beam, a single top load spreading through an inclined compression field to two supports and tied along the soffit (span 2000 mm, depth 1000 mm, thickness 300 mm, $f_c = 30\,\text{MPa}$). Reconstruction accuracy is reported as the root-mean-square error of the minor principal stress σ_2 over the reference field, relative to its root-mean-square magnitude.

5.1 The equilibrium residual is detrimental

A compact ablation on the deep beam at 80×40 gauges, summarised in Table 1, isolates the three loss-balance choices that distinguish the proposed inverse PINN from two natural alternatives. With a dense, noise-free strain field, training on data, compatibility and the boundary conditions reconstructs σ_2 to 6.9%. Adding the strong-form equilibrium residual (4) as a loss term degrades the reconstruction to 23%: the residual, enforced softly, has a spurious low-residual attractor and pulls the solution away from the field that matches the data. Removing the gauge data altogether (the forward-only configuration) leaves the network without a measurement anchor, and the network does not converge to engineering accuracy on this problem in any of the configurations explored in this study. This confirms, on the inverse problem, the mechanism that limits the forward PINN, and it is why Equation (8) omits an equilibrium term.

Table 1: Loss-balance ablation on the deep beam at the dense 80×40 noise-free gauge layout, reporting σ_2 relative-RMS error. "+ equilibrium residual" adds the strong-form PDE residual (4) as a soft loss with the same weighting strategy used in the forward-PINN literature; "forward-only PINN" removes the data term entirely.

configuration	σ_2 rel-RMS
inverse PINN (proposed)	0.069
+ strong-form equilibrium residual	0.230
forward-only PINN (no data)	does not reach engineering accuracy

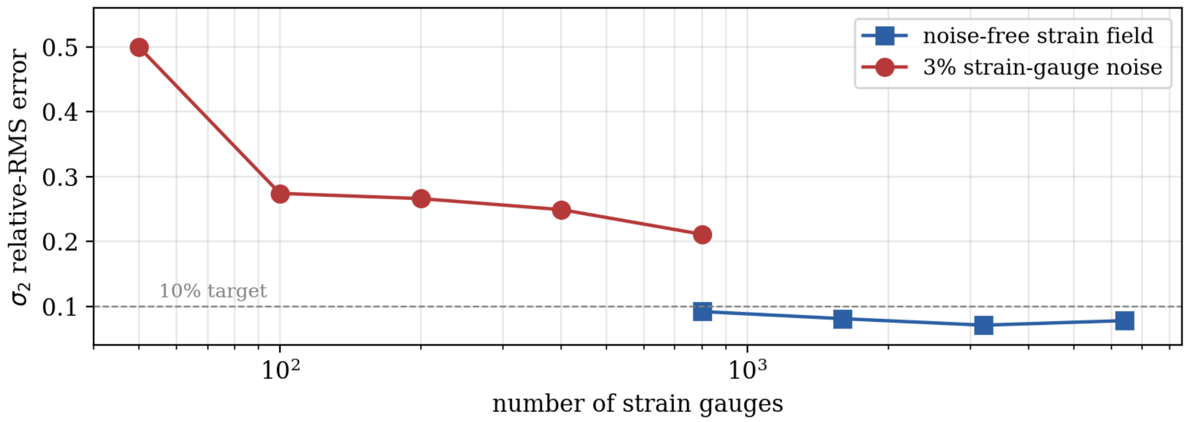

Figure 2: Reconstruction accuracy against measurement density. Without measurement noise the error falls with the number of gauges and then plateaus near the ~7% constitutive floor (Section 5.5). With 3% strain-gauge noise the error is noise-limited at about 20–27% for 100–800 gauges and degrades sharply below about 100 gauges, where the field becomes genuinely under-determined.

5.2 Reconstruction from a dense strain field

With no measurement noise and the strain field sampled densely (an 80×40 gauge grid), the inverse PINN reconstructs the internal stress field to 7.1% in σ_2. The strain components themselves are recovered almost exactly, to 0.8% (ε_x), 1.2% (ε_y) and 1.8% (γ_{xy}); the larger figure in the stress is the constitutive amplification quantified in Section 5.5. The reconstructed field reproduces the inclined compression field and the soffit tie of the deep beam (Fig. 4).

The reconstruction error falls with gauge density and then plateaus: σ_2 is reconstructed to 9.2% at 800 gauges, 8.1% at 1600, 7.1% at 3200 and 7.8% at 6400. The slight uptick at 6 400 gauges is within the stochastic training variability we observe at this noise floor and is not interpreted as a systematic trend; what matters is that beyond a few thousand gauges the strain field is already recovered to within 1–2% and the stress error no longer improves. Neither further densification nor a random-Fourier-feature encoding of the network [22], tested across a range of feature scales, reduces this ~7% floor: it is set by the constitutive amplification (Section 5.5), not by network resolution or data density.

5.3 Measurement density

Holding the noise at a representative gauge level of 3%, the gauge grid was varied from 800 down to about 50 points. The error in σ_2 is about 0.20 at 800 gauges, 0.25–0.27 between 100 and 400 gauges, and 0.50 at 50 gauges. Two observations follow. First, beyond about 100 gauges the reconstruction barely improves with further sampling: at this noise level the accuracy is set by the noise, not by the gauge density. Second, below about 100 gauges the reconstruction degrades sharply, the field becoming genuinely under-determined. Figure 2 shows the two regimes together.

5.4 Measurement noise

With the gauge grid fixed at the dense 40×20 layout, the strain-measurement noise was varied from 0 to 3% (Table 2). The reconstruction stays within 10% in σ_2 for noise up to about 1%, and degrades to about 19% at 3% noise. Noise is therefore the dominant factor in the realistic regime, and the table shows the mechanism: the strain components are recovered to within roughly the measurement noise, while the stress error is several times larger. The increase is not a shortcoming of the network but a propagation effect of the constitutive law, quantified next.

Table 2: Reconstruction error relative to the reference field. The upper block sweeps measurement noise on the deep beam at the 40×20 gauge layout used for the noise study; the lower block reports the two archetypes at their headline dense, noise-free configuration. Values are root-mean-square relative error; σ_2 is the minor principal stress and $\hat{P}$ the load recovered by interior-cut integration (Eq. 10).

archetype	condition	ε_x	ε_y	γ_{xy}	σ_2	$\hat{P}$ error
Deep-beam noise sweep, 40×20 *gauges:*						
deep beam	0 % noise	0.010	0.020	0.021	0.092	n/a
deep beam	1 % noise	0.022	0.058	0.051	0.097	n/a
deep beam	2 % noise	0.043	0.086	0.064	0.136	n/a
deep beam	3 % noise	0.056	0.131	0.080	0.185	n/a
Headline dense, noise-free configuration:						
deep beam	3200 gauges	0.008	0.012	0.018	0.071	0.012
corbel	1500 gauges	0.012	0.011	0.013	0.120	0.083
VK1 wall	2500 gauges	0.021	0.017	0.016	0.015	0.107

5.5 Why strain error is amplified into stress error

The reconstruction error in the stress is consistently several times the error in the strain. This is a property of the cracked CSFM constitutive law, not of the network. A controlled perturbation of the reference field isolates it: scaling the shear strain γ_{xy} by 10% changes the minor principal stress by 34%, an amplification of about 3.4; the corresponding factors are about 1.5 for ε_x and 1.2 for ε_y. The shear strain is amplified most because, through Equation (2), it sets the inclination of the crack and hence of the entire compression field, and the stress is acutely sensitive to that inclination. The consequence for measurement-based assessment is direct: a strain measurement of relative accuracy δ in the shear component yields a CSFM stress field of accuracy no better than roughly $3.4\,\delta$ where the shear strain governs the crack inclination. Reaching a stress field within 10% therefore requires the strain reconstructed within about 3%, which in turn requires both a dense strain field and low measurement noise. The 3.4 factor is geometry-dependent in one important respect: it applies wherever the shear strain controls the principal direction. In purely uniaxial regions, where $\gamma_{xy} \to 0$ and the principal direction is fixed by the axial strains, the amplification collapses toward the axial factors of about 1.5 (ε_x) and 1.2 (ε_y). The 3.4 figure is therefore the relevant bound for typical strut-and-tie D-regions where the shear strain is non-trivial throughout, and a more favourable bound applies in the bending-dominated portions of any specific field (the wall pier of Section 5.8 is an example in which the bending-dominated body and the shear-driven boundary zones contribute differently). Figure 3 summarises the perturbation result.

5.6 Load recovery

The applied load is never supplied to the network. It is recovered after training by integrating the reconstructed vertical stress across a horizontal cut through the interior. With the convention that compression is negative ($\sigma_y < 0$ under the load patch), the load magnitude is

$$\hat{P} = -\int_{x_0}^{x_1} \sigma_y(x,\, y_{\text{cut}})\, t\, dx, \tag{10}$$

where y_{cut} is any horizontal level between the load patch and the supports, (x_0, x_1) is the horizontal extent of the material at that level, and t is the out-of-plane thickness. The minus sign makes $\hat{P}$ positive for a downward applied load. Equilibrium of the material above the cut makes (10) hold identically for any y_{cut} between the load patch and the supports. Cuts in

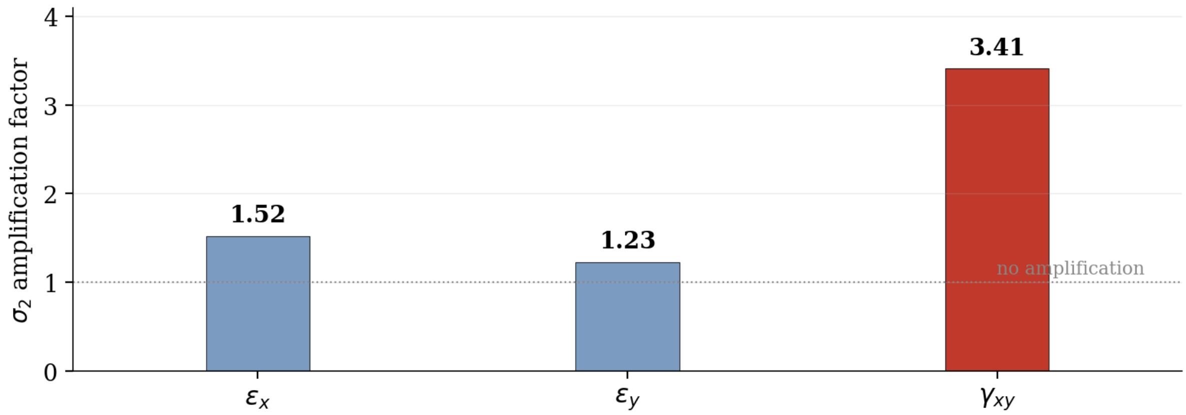

Figure 3: Sensitivity of the minor principal stress σ_2 to a coherent perturbation of each strain component, evaluated on the reference field (perturbations $\varepsilon_x \times 1.10$, $\varepsilon_y \times 1.15$, $\gamma_{xy} \times 1.10$; bar heights are the resulting σ_2 relative-RMS responses divided by the perturbation magnitude). The shear strain γ_{xy} amplifies into the stress by a factor of about 3.4, because it sets the inclination of the crack and hence of the entire compression field; the axial-strain factors are about 1.5 (ε_x) and 1.2 (ε_y). This amplification is the central limitation of any strain-based CSFM stress assessment.

the interior, where the gauge data densely constrain the reconstructed field, give a far more accurate $\hat{P}$ than the boundary integration $\int \sigma_y(x, H)\, t\, dx$ on the load patch itself, where the field is least constrained: the network's spectral bias smooths the discontinuity between the loaded and the unloaded portions of the top edge and the resulting σ_y is biased low.

For the dense, noise-free deep-beam run (10) evaluated at the three cuts $y_{cut} = 0.5H,\ 0.7H,\ 0.85H$ gives $\hat{P} = 777,\ 826,\ 826$ kN, with mean 810 ± 24 kN against the true 800 kN, a 1.2% error. The cut-to-cut spread is itself an honest internal consistency check: it would vanish in a perfectly equilibrated reconstruction and grows as the field deteriorates. The naive boundary estimate for the same run is 543 kN (32% error), illustrating the magnitude of the bias avoided by the interior cut.

An alternative explicit-parameter formulation, treating the load as a scalar $\hat{P}$ co-optimised with the network through the uniform-pressure constraint $\sigma_y = -\hat{P}/(b\,t)$ on the load patch, was tested and proved unworkable: at any constraint weight large enough to pin $\hat{P}$, the load-patch traction dragged the field accuracy, and at smaller weights $\hat{P}$ decoupled from the data. The post-hoc cut-based integration is therefore the recommended route.

5.7 Generalisation to a second archetype: a cantilever corbel

To test that the method extends beyond the simply supported deep beam, the same training pipeline, with no hyper-parameter change, is applied to a cantilever-bracket idealisation of a precast corbel: a 500 mm × 400 mm D-region clamped along its left face (the column attachment, taken as rigid) and loaded over a 100 mm bearing patch at the top of the free end, with $P = 90$ kN, corresponding to a load factor of 1.46 against the continuum-CSFM failure load. The compression flow is qualitatively distinct from the deep beam's: a single inclined strut from the loaded patch down to the bottom corner of the clamped face, balanced by a horizontal top tie. The reinforcement field is correspondingly different, with the 1.2 % tie band placed along the top of the bracket instead of along a soffit.

With 1500 gauges (a 43 × 35 grid, none with knowledge of the load) the reconstruction reaches σ_2 relative-RMS 12.0 %, compared with 7.1 % on the deep beam. The component breakdown is consistent: ε relative-RMS is 0.8–1.3% across the three strain components (deep beam 0.8–1.8%); the stress error is dominated by σ_y and τ_{xy} at ~11 %, again the constitutive amplification in action. The interior-cut load recovery (Eq. 10) gives $\hat{P} = 82.5 \pm 2.5$ kN over three

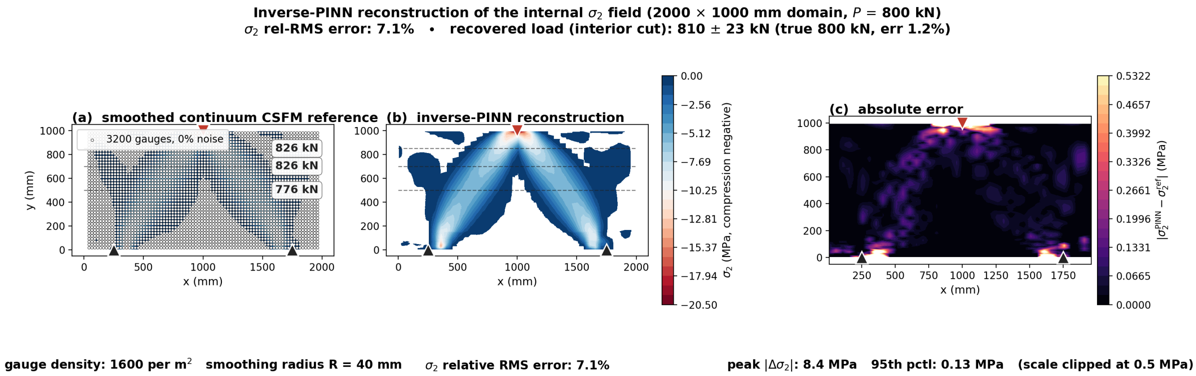

Figure 4: Reconstruction of the deep-beam compression field from 3200 noise-free strain gauges. (a) Smoothed continuum CSFM reference, with the strain-gauge grid overlaid and three dashed interior cuts at $y = 0.5H,\ 0.7H,\ 0.85H$; the recovered load integrated across each cut is annotated. (b) Inverse-PINN reconstruction sharing one σ_2 colour scale with (a), so equal colours read equal MPa. (c) Absolute error $|\sigma_2^{\mathrm{PINN}} - \sigma_2^{\mathrm{ref}}|$ on its own scale, clipped at five times the 95th percentile so the bulk error pattern is visible; the true peak and the clipping value are reported below the panel. The inclined compression field and the soffit tie are recovered; the error concentrates at the load and support patches and along the strut edges. The recovered load quoted in the title is the mean of the three cut integrations (Eq. 10); cut-to-cut spread of ±23 kN sets a self-consistency check.

cuts at $y = 0.25H,\ 0.5H,\ 0.75H$ against the true 90 kN, an 8.3% error and a cut-to-cut spread tighter than the deep beam's. The naive boundary integration, by contrast, gives 66.4 kN (26.2 % error), again confirming the cut method's robustness as a load-recovery diagnostic. Figure 5 shows the corbel reconstruction.

The corbel error is about 1.8 times the deep-beam error. The two geometries differ on several axes at once (the corbel has a smaller domain, a single inclined strut rather than two, and a shorter distance from the data-rich interior to the steepest boundary gradients), so we attribute the gap qualitatively rather than to any single cause. The amplification mechanism (Section 5.5) governs both: relative-RMS strain ~1% in both archetypes, relative-RMS σ_2 several times larger. The result confirms that the inverse-PINN methodology is not specific to the deep-beam geometry; the same loss formulation and hyper-parameters reconstruct the internal stress field of a qualitatively different D-region.

The L-BFGS polish is observed to play a different role on each archetype, in a pattern that may itself be useful as a diagnostic of boundary-traction smoothness. The polish improves the total loss in every case but *degrades* the reconstructed σ_2 in two of the three archetypes (Table 3). The mechanism is clear from the loss components: the second-order step fits the gauges more aggressively, which forces the field to develop steeper gradients near the boundary tractions that the data does not constrain directly. The effect is strongest on the corbel, where the single-strut load path lacks the boundary-traction smoothness of the deep beam's symmetric bearings, mild on the deep beam (the bearings absorb most of the over-fit harmlessly), and marginally beneficial on the wall pier (where the bending-dominated field is smooth throughout). The direction of the L-BFGS effect therefore correlates with how smoothly the boundary tractions are constrained by the geometry, and may itself be read as a flag for geometries where the boundary constraint is insufficient and additional traction-side regularisation would be appropriate.

We mitigate the polish degradation with an *adaptive* L-BFGS that runs in chunks of 100 second-order steps and monitors the compatibility residual: when compat drifts above 1.2× its end-of-Adam value, training stops and the network parameters are restored to the snapshot with the lowest compat seen so far. On the corbel the adaptive variant cuts the polish damage by two-thirds (+0.5 pp instead of +1.5 pp) and is the version reported in Table 3; it leaves the

Table 3: L-BFGS polish effect on σ_2 relative-RMS error per archetype. The reported L-BFGS column uses the adaptive variant described above; the parenthesised values for the deep beam and corbel show what the naive (non-adaptive) L-BFGS gives on the same Adam-trained network. The polish helps only on the smoothest field (VK1) and degrades the other two; the adaptive variant cuts the corbel's degradation by two-thirds.

archetype	after Adam	after adaptive L-BFGS	$\Delta\sigma_2$
deep beam (3200 gauges)	0.064	0.071 (naive: 0.071)	+0.007 (worse)
corbel (1500 gauges)	0.115	0.120 (naive: 0.130)	+0.005 (worse)
VK1 wall (2500 gauges)	0.018	0.015 (naive: 0.015)	−0.003 (better)

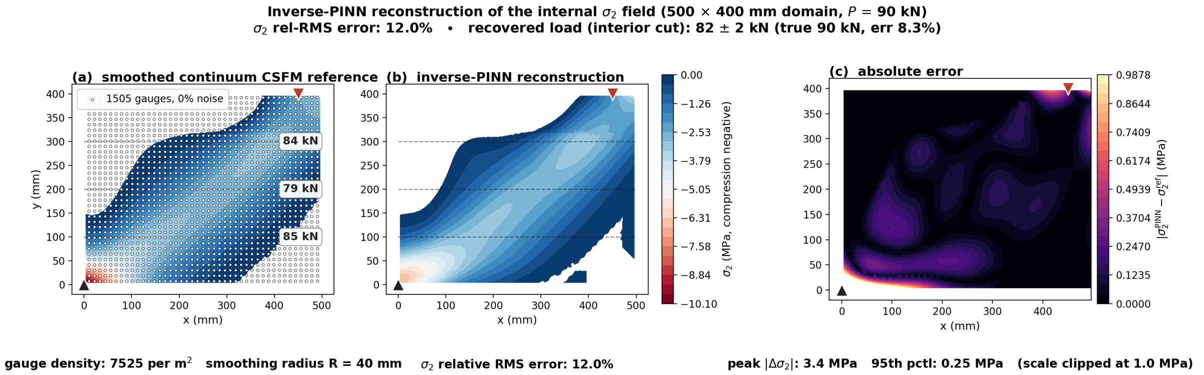

Figure 5: Reconstruction of the corbel compression field from 1500 noise-free strain gauges. Same panel layout and conventions as Figure 4. The single inclined strut from the loaded patch to the bottom of the clamped face, and the horizontal top tie, are both recovered; error concentrates near the loaded patch and the clamped corner. Cut-line integration at $y \in \{0.25H, 0.5H, 0.75H\}$ recovers $\hat{P} = 82\,\text{kN}$ against the true 90 kN.

deep beam and VK1 results unchanged because their L-BFGS trajectories do not trigger the rollback criterion. The residual half-percentage-point degradation on the corbel reflects an intrinsic limitation of the monitor: the compatibility residual quantifies the agreement between the strain network and the gradient of the displacement network, not the agreement between either network and the true field; the two networks can grow mutually more consistent (lower compat) while jointly drifting from the reference. A direct σ_2 proxy would require either a held-out set of gauges or an alternative physics-based criterion (e.g. the free-edge traction residual), neither of which is explored here. The practical recommendation is therefore to treat the L-BFGS polish as a diagnostic rather than as a default: train with Adam alone on geometries with non-smooth boundary tractions, and only apply the adaptive polish when its effect on a reference-equipped test problem is verified.

5.8 Application to the Bimschas wall pier (VK1)

To bridge from synthetic benchmarks to a D-region representative of an in-service tested structure, the method is applied to the geometry, materials, reinforcement and loading of specimen *VK1* of Bimschas [3], included in the experimental validation suite of the CSFM (Kaufmann et al. [9], Section 6.3). VK1 is a 1500 mm × 3700 mm × 200 mm cantilever wall-type bridge pier with vertical (flexural) reinforcement $\rho_l = 0.82\,\%$ and horizontal (shear) reinforcement $\rho_t = 0.08\,\%$, fixed at its base and loaded by a constant axial compression $N = 1370\,\text{kN}$ and a horizontal force V applied at an effective height $H_{\text{eff}} = 3300\,\text{mm}$. The materials are $f_c = 35\,\text{MPa}$ and $f_y = 515\,\text{MPa}$. The experimentally measured ultimate horizontal force was $V_{u,\text{exp}} = 725\,\text{kN}$ (failure mode: concrete crushing with flexural yield, Bimschas 3).

We apply the inverse-PINN pipeline at $V = 300\,\text{kN}$, corresponding to a service-load level of 0.41 $V_{u,exp}$. The choice is conservative: our continuum-CSFM solver, which uses a first-element compression failure check rather than the integrated plastic-redistribution CSFM implementation used in Kaufmann et al. [9], converges with a load-factor reserve of 1.95 at this level. (For reference, the book's CSFM analysis predicts $V_{u,calc} = 744$ kN, within ~3% of the measurement; reproducing the ultimate state in the inverse PINN would require a solver that captures the full plastic redistribution and is left to future work.)

With 2500 noise-free synthetic gauges, the inverse PINN reconstructs the internal compression field to σ_2 relative-RMS 1.5 %, the smallest reconstruction error of any archetype in this study. The strain components are recovered to 0.021, 0.017 and 0.016 relative-RMS in ε_x, ε_y and γ_{xy} respectively, and the peak compression magnitude matches the continuum reference to within 6 % (31.5 MPa versus 33.6 MPa).

The interior-cut recovery is applied to *both* unknowns. Three horizontal cuts at $y \in \{0.25H_{eff}, 0.45H_{eff}, 0.70H_{eff}\}$ give:

$$\hat{N} = \tfrac{1}{3}\sum_{y}\int_0^L (-\sigma_y)\, t\, dx = 1474(66)\,\text{kN} \quad (\text{true } 1370\,\text{kN},\ \text{error } 7.6\,\%), \tag{11}$$

$$\hat{V} = \tfrac{1}{3}\sum_{y}\int_0^L \tau_{xy}\, t\, dx = 332(7)\,\text{kN} \quad (\text{true } 300\,\text{kN},\ \text{error } 10.7\,\%). \tag{12}$$

The cut-to-cut spread of ±7 kN on V (~2% of the mean) is the tightest of any archetype, consistent with the smooth, bending-dominated nature of the field. The reconstruction therefore recovers both unknown forces simultaneously from a single trained network, with the axial load and the shear load each derived from the appropriate stress component of the same reconstructed field.

This application anchors the method to a documented tested specimen without overclaiming: the gauge data are still synthetic, sampled from a continuum-CSFM forward solution; the geometry, reinforcement, material and the value $N = 1370$ kN are taken directly from the test. The headline link to experiment is the $V_{u,exp} = 725$ kN context value, which sets the scale of the analysis. A full application to instrumented test data, with strain measurements from digital image correlation or distributed fibre-optic sensing replacing the synthetic gauges, is the natural next step and is reserved for the planned companion study.

Figure 6 shows the reconstructed compression field.

6 Discussion

The results place the inverse PINN, and physics-informed learning for the CSFM generally, on a clear footing. Three points stand out.

6.1 The inverse formulation suits the PINN setting

A forward PINN trained only against the governing physics, with no measurement anchor, did not reach engineering accuracy in any of the configurations explored in this study. The strong-form equilibrium residual admits a family of near-equilibrated fields, and training tends to converge to a diffuse one rather than to the true solution. The inverse formulation does not appear to share this difficulty because the measurements anchor the solution. With the data in place, the equilibrium residual was found to be harmful as a soft penalty and is not required once compatibility, the constitutive law and the traction conditions are imposed. For this problem the data-assimilating network appears to be the more useful object.

6.2 The constitutive law sets the achievable accuracy

The amplification quantified in Section 5.5 appears to be a primary limitation on the achievable accuracy, and the mechanism is physical rather than numerical. Because the cracked CSFM

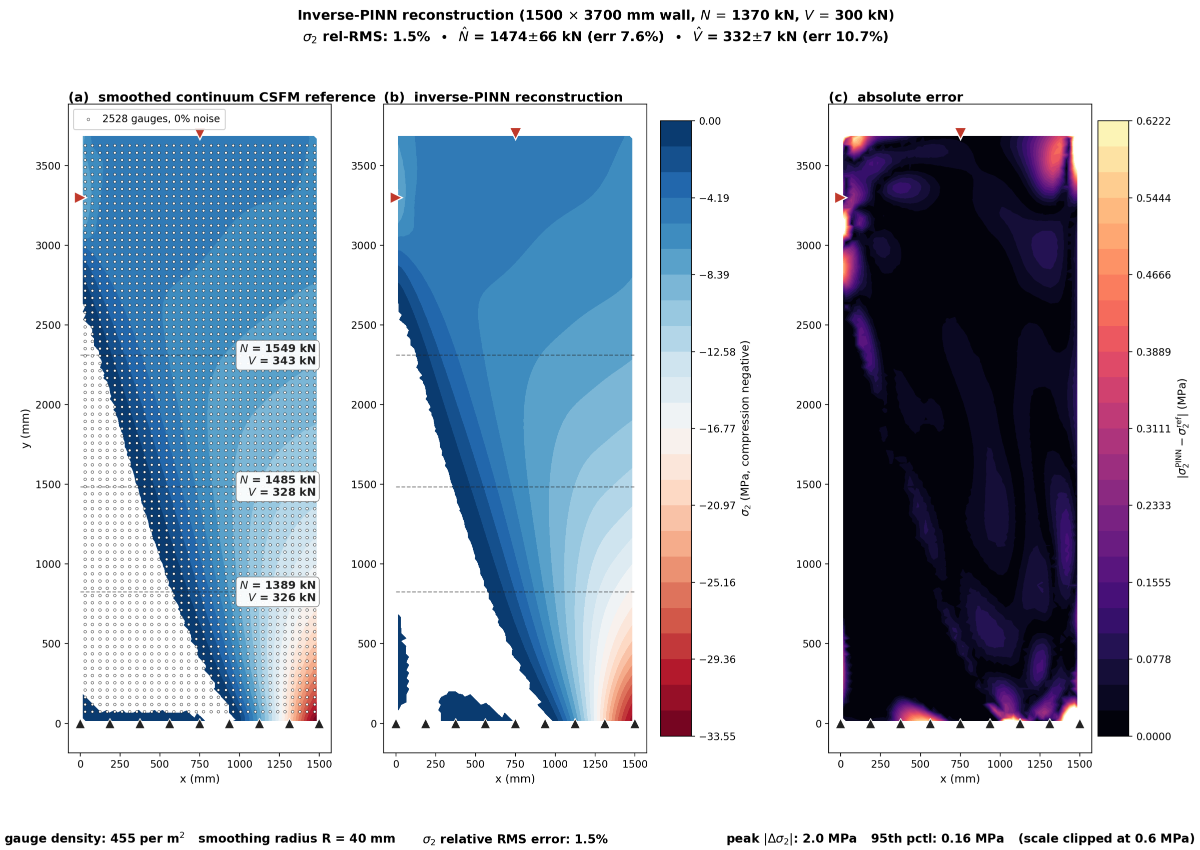

Figure 6: Reconstruction of the VK1 wall pier compression field at 300 kN horizontal force (41 % of $V_{u,exp}$) and 1370 kN axial compression, from 2500 noise-free synthetic strain gauges. Same panel layout and conventions as Figure 4. The cut annotations on the reference panel report the axial force integrated across each horizontal interior cut, mean $\hat{N} = 1474(66)$ kN versus the applied 1370 kN; the shear force recovered from the same cuts by integrating τ_{xy} instead gives $\hat{V} = 332(7)$ kN versus the applied 300 kN. The σ_2 relative-RMS error is 1.5 %, the smallest of any archetype in this work, consistent with the smooth, bending-dominated nature of the wall field.

stress depends on the crack inclination, and the inclination on the shear strain, a stress-field reconstruction inherits roughly 3.4 times the relative error of the strain it is built from. This result may be useful for measurement-based assessment more broadly: it suggests how accurate, and how dense, a strain measurement must be before a CSFM stress check built on it can be trusted, and it likely applies to any method that maps measured strain to CSFM stress, not only to the network of this study.

6.3 What the method delivers

Within these limits the inverse PINN occupies a niche that a conventional forward finite-element analysis does not address: it provides a complete, mesh-free, continuous CSFM stress field, and an estimate of the load, from strain measurements and the physics alone. With a dense strain field at low measurement noise, of the kind a full-field optical measurement can provide, the reconstruction in the experiments reported here is within about 7% on the deep beam; with sparser or noisier data the error grows in a manner consistent with the amplification result. The method may therefore be useful as a tool for converting a measured strain field into a CSFM stress assessment, accompanied by an explicit error budget.

7 Limitations and future work

The study uses three plane-stress archetypes: a simply supported deep beam, a cantilever-bracket corbel, and the geometry of the VK1 wall pier of Bimschas [3]. The VK1 application is anchored to a tested specimen but the gauge data are still synthetic, sampled from a continuum-CSFM forward solution rather than from physical measurements; replacing them with real digital-image-correlation or distributed-strain-sensing data from an instrumented specimen is the natural next step and the focus of a planned companion study. The VK1 application is also limited to service-load conditions (~41 % of the measured ultimate) because our continuum-CSFM solver implements a first-element compression failure check rather than the plastic-redistribution CSFM implementation of Kaufmann et al. [9]; spanning to the ultimate state, where the inverse PINN could be compared directly to $V_{u,exp}$, requires a solver with the full integrated ULS treatment. The remaining companion D-regions of the CSFM (pile caps, hammerheads, dapped beams, joints) and the three-dimensional case, in which the interior is genuinely hidden from surface measurement and the inverse framing is strongest, are also left to future work. We did attempt an L-shaped corbel geometry, with the column explicitly modelled below the bracket arm: that case proved hostile to the current PINN architecture because the rectangular bounding box has a large empty quadrant in which the network's high-frequency degrees of freedom oscillate unconstrained and then propagate into the concrete via spatial continuity. A domain-masked network architecture would address this; the present results therefore apply to D-regions whose bounding box is close to fully filled. The reconstruction accuracy is governed by the constitutive amplification and by measurement noise: sub-10% stress accuracy is reached for a densely sampled strain field at measurement noise up to about 1%, and degrades for sparser or noisier data as quantified above. The reference solution is itself a finite-element field recovered by smoothing, with a residual mesh uncertainty of about three per cent in σ_2; the measured 7% reconstruction error is therefore an upper bound on the network's true error against the converged CSFM field. Subtracting the reference uncertainty in quadrature gives a network error of approximately 6.3%, and further refining the reference would tighten this bound without altering the regime of the result.

8 Conclusions

This study formulated the recovery of the internal stress field of a concrete D-region as an inverse problem and solved it with a physics-informed neural network that assimilates strain measurements with the Compatible Stress Field Method. The architecture pairs a strain network and a displacement network coupled by kinematic compatibility, and embeds the cracked-membrane constitutive law as a differentiable map so that the reconstructed stress is, by construction, a CSFM stress field.

On a simply supported deep beam the network reconstructs the internal σ_2 field to within 7.1% of an independent continuum CSFM solution from a dense, low-noise strain field, and the error grows with measurement noise broadly in line with the constitutive amplification result. Applying the same pipeline, with no hyper-parameter change, to a cantilever-bracket corbel, a qualitatively distinct D-region whose load path is a single inclined strut closed by a top tie, yields 12.0% σ_2 accuracy. Applied to the geometry of an experimentally tested cantilever wall pier from the validation suite of the CSFM (specimen VK1 of Bimschas [3], included in Kaufmann et al. [9]), the same pipeline reconstructs the field to 1.5% at service-load conditions and simultaneously recovers both the horizontal load and the axial compression by integrating the appropriate stress components across the same interior cuts. The method therefore extends beyond synthetic benchmarks to D-region geometries representative of tested specimens; the bending-dominated smoothness of the wall field makes it easier for the network than the sharp strut-and-tie patterns of the deep beam and corbel.

The applied load, never supplied to the network, follows from integrating the reconstructed vertical stress across an interior horizontal cut, where the gauge data densely constrains the field. That integral recovers P to within 1.2% on the deep beam and 8.3% on the corbel, and the cut-to-cut spread doubles as a self-consistency check. On the wall pier the procedure generalises naturally to simultaneous recovery of both unknown forces: integrating the reconstructed σ_y across an interior horizontal cut returns the axial compression to 7.6%, and integrating τ_{xy} across the same cuts returns the horizontal shear to 10.7%. Integrating the same stresses across the load patch on the boundary itself, by contrast, gives substantially larger errors (32% and 26% respectively on the deep beam and corbel), because the network's spectral bias smooths the discontinuity between the loaded and unloaded portions of the top edge. The interior cut therefore appears to be the more reliable route to load recovery from a strain-based reconstruction, at least for the geometries considered here. An alternative explicit co-optimisation of the load through a uniform-pressure constraint was tested in this study and proved unworkable.

Two methodological findings underpin these results. The strong-form equilibrium residual proves detrimental as a training objective, and a purely forward physics-only network is not competitive: the data-assimilating formulation, coupled to compatibility and the traction conditions through the embedded constitutive map, suffices, and equilibrium emerges from the reconstruction rather than constrains it. The cracked CSFM constitutive law also appears to impose an error budget that bounds any strain-based stress reconstruction: a relative-RMS strain error δ propagates into a σ_2 error of roughly $3.4\,\delta$, because the crack direction is governed by the shear strain and that direction sets the compression field. The amplification factor is expected to apply to any strain-based CSFM stress assessment, not only to the network of the present work, and is a central practical result of this study.

Taken together, the three reconstructed archetypes and the load-recovery procedure define a usable workflow for measurement-based CSFM assessment: gauge strain in, full σ_2 field and applied load out, with a transparent and quantified error budget governed by the constitutive amplification factor.

Data and code availability

The differentiable constitutive port, the inverse-PINN training code and the trained models will be released in a public repository and archived with a permanent identifier on publication.

Declaration of competing interest

The authors declare no competing financial interests or personal relationships that could have influenced the work reported in this study.

References

[1] Vera M. Balmer, Walter Kaufmann, and Michael A. Kraus. Physics-informed neural networks for nonlinear analysis of reinforced concrete beams. In *Lecture Notes in Civil Engineering*, volume 494. Springer, 2024. doi:10.1007/978-3-031-60271-9_24.

[2] Jan-Hendrik Bastek and Dennis M. Kochmann. Physics-informed neural networks for shell structures. *European Journal of Mechanics - A/Solids*, 97:104849, 2023.

[3] Martin Bimschas. *Displacement-Based Seismic Assessment of Existing Bridges in Regions of Moderate Seismicity*. PhD thesis, ETH Zurich, 2010. Doctoral dissertation, ETH Zurich, Switzerland.

[4] Rafael Bischof and Michael A. Kraus. Multi-objective loss balancing for physics-informed deep learning. *Computer Methods in Applied Mechanics and Engineering*, 439:117914, 2025. doi: 10.1016/j.cma.2025.117914.

[5] Branko Glisic and Daniele Inaudi. *Fibre Optic Methods for Structural Health Monitoring*. Wiley, 2007.

[6] Ehsan Haghighat, Maziar Raissi, Adrian Moure, Hector Gomez, and Ruben Juanes. A physics-informed deep learning framework for inversion and surrogate modeling in solid mechanics. *Computer Methods in Applied Mechanics and Engineering*, 379:113741, 2021.

[7] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, Paris Perdikaris, Sifan Wang, and Liu Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3(6):422–440, 2021.

[8] Walter Kaufmann and Peter Marti. Structural concrete: Cracked membrane model. *Journal of Structural Engineering*, 124(12):1467–1475, 1998.

[9] Walter Kaufmann, Jaime Mata-Falcon, Manuel Weber, and Tena Galkovski. *Compatible Stress Field Design of Structural Concrete*. ETH Zurich, 2020.

[10] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015.

[11] Sandesh Lamsal. DeepForm: A graph neural network surrogate for real-time tensile membrane form-finding across anticlastic typologies. engrXiv preprint: https://doi.org/10.31224/7269, 2026.

[12] Sandesh Lamsal and Rubi Bhandari. A strut and tie neural network surrogate for failure-load prediction of concrete d-regions designed by the compatible stress field method. engrXiv preprint: https://doi.org/10.31224/7250, 2026.

[13] Dong C. Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(1):503–528, 1989.

[14] Peter Marti. Basic tools of reinforced concrete beam design. *ACI Journal*, 82(1):46–56, 1985.

[15] Aurelio Muttoni, Joseph Schwartz, and Bruno Thurlimann. *Design of Concrete Structures with Stress Fields*. Birkhauser, 1997.

[16] Vien Minh Nguyen-Thanh, Xiaoying Zhuang, and Timon Rabczuk. A deep energy method for finite deformation hyperelasticity. *European Journal of Mechanics A/Solids*, 80:103874, 2020.

[17] Fabrice Pierron and Michel Grédiac. *The Virtual Fields Method: Extracting Constitutive Mechanical Parameters from Full-Field Deformation Measurements*. Springer, 2012.

[18] Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.

[19] Prajit Ramachandran, Barret Zoph, and Quoc V. Le. Searching for activation functions. *arXiv preprint*, 2017. arXiv:1710.05941.

[20] Esteban Samaniego, Cosmin Anitescu, Somdatta Goswami, Vien Minh Nguyen-Thanh, Hongwei Guo, Khader Hamdia, Xiaoying Zhuang, and Timon Rabczuk. An energy approach to the solution of partial differential equations in computational mechanics via machine learning: Concepts, implementation and applications. *Computer Methods in Applied Mechanics and Engineering*, 362:112790, 2020.

[21] Jorg Schlaich, Kurt Schafer, and Mattias Jennewein. Toward a consistent design of structural concrete. *PCI Journal*, 32(3):74–150, 1987.

[22] Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. In *Advances in Neural Information Processing Systems*, volume 33, 2020.

[23] Frank J. Vecchio and Michael P. Collins. The modified compression-field theory for reinforced concrete elements subjected to shear. *ACI Journal*, 83(2):219–231, 1986.

[24] Sifan Wang, Yujun Teng, and Paris Perdikaris. Understanding and mitigating gradient flow pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 43 (5):A3055–A3081, 2021.