

Identifiability and experimental design for inverse physics-informed engineering models

Luis Torres

Abstract

An inverse physics-informed neural network can minimize its objective accurately while returning nonunique or weakly constrained material and operating parameters. This review explains why optimization success must be separated from parameter identifiability. Exact examples from heat conduction, pressure-driven flow, and axial deformation show how observations can determine a parameter combination without determining its individual components. Local sensitivity matrices, dimensionless scaling, nuisance parameters, and profile objectives are then used to organize practical diagnostics. Published flow and viscoplastic PINN studies provide context for representation and constitutive modeling, while inverse-problem analysis supplies qualified stability results. A proposed experimental-design workflow links sensor placement and loading choices to the engineering parameter of interest.

1. Fitting a field and identifying a parameter are different tasks

In a forward problem, material parameters, loads, and boundary conditions determine a physical field. In an inverse problem, selected observations are used to infer some of those inputs. PINNs can combine an observation loss with governing-equation residuals and optimize network weights together with unknown parameters [1, 2]. However, a differential equation does not guarantee that the observations distinguish all candidate parameters. If two parameter choices generate exactly the same measured response, no optimizer can distinguish them without additional information.

Structural identifiability concerns uniqueness under idealized, sufficiently informative observations within a specified model. Practical identifiability concerns what finite, noisy, incomplete measurements can resolve. A parameter can be structurally identifiable but practically uncertain because its effect is too small, strongly correlated with another effect, or masked by uncertain boundary conditions. Reviews of structural identifiability and observability provide a useful conceptual framework, although results developed for nonlinear dynamical systems do not automatically apply to every PDE inverse problem [3].

This distinction is especially important in engineering calibration. A network may fit displacement measurements while compensating an incorrect modulus with an incorrect applied load. It may fit temperature while inferring an effective source strength rather than a conductivity. It may reproduce a velocity profile while identifying only a pressure-gradient-to-viscosity ratio. A credible inverse study must state exactly which quantities are known, inferred, measured, or regularized.

2. Three exact engineering examples

Consider steady one-dimensional conduction in a homogeneous rod with positive constant conductivity k , no internal source, and prescribed endpoint temperatures T_o and T_L . The governing equation and solution are

$$\begin{aligned} -k T'' &= 0, T(0) = T_o, T(L) = T_L. \\ T(x) &= T_o + (T_L - T_o)x/L. \end{aligned}$$

Every positive k produces the same temperature field. Even exact temperatures at every interior point cannot identify k . A PINN can satisfy the PDE and all temperature data for an entire continuum of conductivity values. A conductivity penalty or initialization preference may select one value, but the selected value is not identified by these observations. Measuring the heat flux adds information because $q = -k(T_L - T_o)/L$, provided the temperature difference is nonzero and the flux is known accurately.

With a uniform internal source Q and prescribed endpoint temperatures, the curvature satisfies $T'' = -Q/k$. If Q is known and nonzero, ideal curvature data can determine k . If both Q and k are unknown, scaling them by the same positive factor leaves the entire temperature solution unchanged. The observations determine Q/k . This simple invariance is an exact nonidentifiability result, not merely a correlation observed during numerical training.

For transient conduction with volumetric heat capacity C , the equation $C \partial T / \partial t = k \partial^2 T / \partial x^2$ depends on thermal diffusivity $\alpha = k/C$ when there are no additional flux or source terms introducing separate information. Temperature histories can identify α under appropriate excitation and observations, but cannot generally separate k from C in that formulation. A known heat-flux boundary condition can alter this conclusion because its dimensional relation includes k explicitly. Identifiability is therefore a property of the complete experiment, not only the interior PDE.

Next consider fully developed laminar flow between stationary parallel plates at $y = \pm h$. With pressure-gradient magnitude $G = -dp/dx > 0$ and viscosity μ , the velocity is

$$\begin{aligned} u(y) &= [G/(2\mu)](h^2 - y^2). \\ u(y; cG, c\mu) &= u(y; G, \mu), c > 0. \end{aligned}$$

A complete noiseless velocity profile determines G/μ when h is known. It does not identify G and μ separately. An independently measured pressure drop over a known length can resolve G and thereby μ . Uncertain plate spacing introduces a further geometric effect because the amplitude contains h^2 and the observed coordinate mapping depends on h . This is directly relevant to flow reconstruction, where measurement geometry and boundary information matter as much as network expressiveness [4, 5, 6].

Finally, a uniform elastic bar under known axial force F has elongation $\delta = FL/(EA)$. If length L is known but modulus E and cross-sectional area A are both unknown, elongation identifies their

product EA . If the force is also uncertain, only a ratio involving F is constrained. Solid-mechanics PINNs may enforce equilibrium accurately while retaining this ambiguity [7]. Adding a direct stress output changes the numerical formulation, but does not supply an independent force or area measurement [8].

Experiment	Combination resolved by ideal observations	Additional information that can separate parameters
Source-free steady conduction, prescribed temperatures	No conductivity information from temperature alone	Nonzero measured heat flux
Steady conduction with unknown uniform source	Q/k	Independently known source or conductivity
Transient conduction under the stated temperature conditions	k/C	Independent heat capacity or suitable flux information
Parallel-plate velocity profile	G/μ	Pressure drop and known channel dimensions
Axial elongation of a bar	EA when force is known	Independently measured area or modulus

3. Sensitivity, rank, and the limits of local analysis

Let θ denote physical parameters and let $m(\theta)$ be the predicted measurement vector after solving the forward problem. For observations y with a parameter-independent positive-definite noise covariance Σ , a local Gaussian least-squares analysis uses the sensitivity matrix S and information matrix I :

$$S_{ij} = \partial m_i / \partial \theta_j, I = S^T \Sigma^{-1} S.$$

$$\delta m \approx S \delta \theta.$$

If a nonzero perturbation $\delta\theta$ lies in the null space of S , that direction has no first-order effect on the observations. Small singular values indicate weak local sensitivity after scaling. For an exact scaling invariance, as in the plate-flow example, the deficient direction corresponds to a genuine continuum of equivalent parameters. Full local rank, however, does not prove global uniqueness. Nonlinear maps can have distinct separated parameter values that generate the same observations.

Parameter scaling matters when interpreting singular values. A sensitivity with respect to conductivity in one unit system cannot be compared directly with a sensitivity to a dimensionless coefficient. For positive parameters, logarithmic coordinates give sensitivities to relative changes. Alternatively, normalize each parameter by a justified engineering scale. The rank of an exactly transformed Jacobian is preserved by invertible scaling, but numerical conditioning and threshold-based decisions can change substantially.

The derivatives must also correspond to the physical inverse problem. Differentiating a neural output with respect to a parameter while freezing network weights can be inappropriate if those weights were separately optimized for each parameter value. One needs either a reliable differentiable forward surrogate over parameters, an implicit sensitivity of the governing solve, or an explicitly reoptimized profile. The joint Hessian of all network weights and parameters is not automatically an information matrix for physical parameters because neural reparameterizations introduce additional degeneracies.

Conditional stability analyses for selected PINN inverse problems help connect PDE residuals, observation errors, and solution errors [9]. Their results address specified PDE and observation settings. They do not establish that arbitrary material parameters in nonlinear plasticity are identifiable. Forward convergence results likewise do not supply inverse uniqueness [10]. This separation prevents an otherwise sound approximation theorem from being used to justify an unsupported parameter-recovery claim.

4. Nuisance parameters and practical uncertainty

Unknown boundary conditions, sensor offsets, source locations, geometry tolerances, and time synchronization are nuisance parameters when the primary target is a material property. Fixing them to inaccurate nominal values can make the reported material estimate appear precise while introducing bias. A sensitivity analysis should therefore include plausible nuisance directions. If conductivity and boundary heat-transfer coefficient have similar effects over the observed time window, holding the coefficient fixed can conceal the true uncertainty in conductivity.

With target parameters θ and nuisance parameters η , partition the local information matrix into blocks. If the nuisance block is invertible, the effective local information after accounting for nuisance variation is the Schur complement:

$$I_{\text{effective}} = I_{\theta\theta} - I_{\theta\eta} I_{\eta\eta}^{-1} I_{\eta\theta}.$$

The subtracted term is positive semidefinite under the stated information-matrix assumptions. Thus treating nuisance parameters as unknown cannot increase this local information relative to fixing them exactly. This is a local Gaussian approximation, not a general formula for nonlinear posterior uncertainty. It nevertheless explains why an apparently well-conditioned fixed-boundary calibration can become uncertain when realistic boundary errors are admitted.

Bayesian PINNs offer a probabilistic framework for uncertain fields and parameters [11]. A concentrated posterior can still reflect a strong prior rather than informative observations. Report prior and posterior ranges, sensitivity to prior choices, and predictive checks on held-out experiments. Ensembles expose some dependence on initialization and fitting choices, but ensemble spread alone need not represent structural ambiguity or measurement-model error [12]. Neither approach removes the need for an identifiable experiment.

Measurement correlation also matters. Hundreds of neighboring pixels or consecutive time samples can share systematic errors. Treating them as independent with a diagonal covariance

can overstate information. Separate repeatability noise from calibration error, and report how spatial and temporal correlations are modeled. A validation experiment with a different sensor calibration or loading condition can be more informative than adding many nearly redundant points to the same acquisition.

5. What neural architecture and physics penalties can change

Frequency-aware architectures can represent oscillatory fields and parameter-dependent flow patterns more effectively in suitable problems [13, 14, 15, 16]. This reduces one possible source of forward approximation error. It cannot resolve an exact invariance in the observation map. Zhang et al.'s geometry-conditioned flow predictions demonstrate a forward-surrogate capability; they are not a proof that geometry, viscosity, and inlet conditions can all be inferred uniquely from an arbitrary sparse measurement set [13].

Arora et al. model rate- and temperature-dependent viscoplastic responses with appropriate mechanical state outputs and scaled losses [17]. For inverse use, additional questions arise: whether the loading paths distinguish elastic, viscous, and hardening effects; whether temperature is independently measured; and whether internal-state uncertainty is represented. Thermodynamic neural formulations can constrain constitutive structure, but their admissible parameter families still require informative experiments [18]. A monotonic test may fail to separate parameters that unloading or rate jumps could distinguish.

Optimization failure can mimic nonidentifiability, while nonidentifiability can mimic optimization failure. Gradient imbalance, neural-tangent-kernel effects, and difficult PDE regimes can create flat or badly conditioned training dynamics even for an identifiable physical problem [19, 20, 21]. Adaptive weights, causal training, and gradient-enhanced methods may improve optimization, but they do not change the information content of a fixed dataset [22, 23, 24, 25]. A useful diagnostic is to solve a simplified forward or inverse reference problem with a conventional method before attributing all uncertainty to the PINN.

Similarly, residual-adaptive and importance sampling allocate physical collocation points [26, 27]. They can reduce approximation error where the field is poorly resolved. They do not create new experimental measurements of an unknown parameter. Sparse-sensor methods address a different design variable: where to acquire independent observations [28]. Keeping these two point sets separate avoids claiming that more collocation points are equivalent to a more informative experiment.

6. A reproducible identification study

Begin by writing the complete forward model and measurement operator, including sensor averaging, coordinate transformations, and unknown boundary quantities. Search analytically for transformations that preserve predictions, as in the conductivity-source and viscosity-pressure examples. Then compute scaled sensitivities around several plausible parameter values using a converged reference model. A single well-conditioned Jacobian at one nominal point is insufficient if the relevant parameter range contains insensitive or ambiguous regions.

Next form a profile objective for each important physical parameter. Fix that parameter at a sequence of values and reoptimize the other physical parameters, nuisance variables, and necessary field representation. A flat profile indicates that substantially different values fit similarly; multiple separated minima suggest possible nonuniqueness or optimization difficulties. Repeat the profile with alternative initializations and improved forward accuracy to distinguish numerical artifacts from persistent ambiguity. Confidence thresholds require a stated noise model and suitable statistical assumptions; arbitrary training-loss increments should not be presented as universal confidence intervals.

Use synthetic recovery tests in stages. First employ a conventional forward solver with known parameters and assess noiseless recovery. Then add realistic noise, nuisance variation, sparse sensing, and finally model mismatch. Avoid generating all synthetic observations with the identical discretization and approximation used in inversion, which can produce an unrealistically favorable test. Thermal and additive-manufacturing settings are useful examples because source shape, boundary loss, and material properties can all compensate each other [29, 30].

Report parameter error where truth is known, predictive error on unused experiments, interval coverage across repeated synthetic datasets, and sensitivity to model discrepancy. Compare with a conventional inverse solver and include total computational cost. Forward comparisons show that PINNs are not automatically the most accurate or economical numerical choice [31]. Benchmarks and reusable frameworks can aid reproducibility, but their forward solution scores do not establish inverse identifiability [32, 33, 34]. Training recommendations should be documented so numerical choices can be separated from physical information [35].

7. Designing a more informative engineering experiment

Select experiments that break the dominant ambiguity. For the steady conduction example, a measured nonzero flux is more decisive than another interior temperature. For pressure-driven flow, add a pressure-drop measurement rather than densely duplicating the same velocity profile. For a bar, measure area and force independently before interpreting elongation as a modulus measurement. In viscoplastic calibration, vary loading rate, temperature, and path in ways that produce distinguishable state trajectories, subject to the validity of the constitutive model.

Local optimal-design criteria summarize different goals. Maximizing the determinant of a scaled information matrix seeks a small local uncertainty volume; maximizing its smallest eigenvalue

targets the weakest direction; minimizing the variance of one quantity targets a specific engineering decision. These criteria depend on the assumed parameter point, noise covariance, and candidate experiments. A robust design evaluates several plausible parameter settings and includes practical constraints such as sensor access, acquisition cost, and thermal or mechanical operating limits.

Repeated excitation should alter the physics, not merely repeat an uninformative condition. Multiple steady conduction tests with prescribed temperatures and zero source still do not identify conductivity from temperature alone. Increasing the number of samples does not repair the invariance. Conversely, a carefully chosen transient or flux-controlled test can introduce the missing parameter dependence. Analytical structure should guide data acquisition before expensive architectural searches.

8. Conclusions

Inverse PINNs inherit the identifiability limits of the underlying experiment. Exact PDE enforcement, sophisticated features, and extensive collocation cannot distinguish parameter combinations that produce identical observations. The engineering remedy is to identify invariances, quantify scaled sensitivities and nuisance effects, and design measurements that break the remaining ambiguity.

Reliable parameter estimation therefore requires two separate demonstrations: that the numerical method solves the stated inverse problem accurately, and that the stated inverse problem contains enough information for the claimed parameters. Reporting both makes fitted values scientifically interpretable and useful for subsequent engineering decisions.

References

- [1] Maziar Raissi, Paris Perdikaris, 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). DOI.
- [2] George Em Karniadakis, Ioannis G. Kevrekidis, Lu Lu, et al.. Physics-informed machine learning. *Nature Reviews Physics* 3, 422-440 (2021). DOI.
- [3] Alejandro F. Villaverde. Observability and Structural Identifiability of Nonlinear Biological Systems. arXiv:1812.04525 (2018). arXiv.
- [4] Chengping Rao, Hao Sun, Yang Liu. Physics-informed deep learning for incompressible laminar flows. arXiv:2002.10558 (2020). arXiv.
- [5] Maziar Raissi, Alireza Yazdani, George Em Karniadakis. Hidden Fluid Mechanics: A Navier-Stokes Informed Deep Learning Framework for Assimilating Flow Visualization Data. arXiv:1808.04327 (2018). arXiv.
- [6] Maziar Raissi, Zhicheng Wang, Michael S. Triantafyllou, George Em Karniadakis. Deep Learning of Vortex Induced Vibrations. arXiv:1808.08952 (2018). arXiv.
- [7] Ehsan Haghighat, Maziar Raissi, Adrian Moure, et al.. A deep learning framework for solution and discovery in solid mechanics. arXiv:2003.02751 (2020). arXiv.
- [8] R. Sharma, Y. B. Guo. Thermal-Mechanical Physics Informed Deep Learning For Fast Prediction of Thermal Stress Evolution in Laser Metal Deposition. arXiv:2412.18786 (2024). arXiv.

- [9] Siddhartha Mishra, Roberto Molinaro. Estimates on the generalization error of Physics Informed Neural Networks (PINNs) for approximating a class of inverse problems for PDEs. arXiv:2007.01138 (2020). arXiv.
- [10] Yeonjong Shin, Jerome Darbon, George Em Karniadakis. On the convergence of physics informed neural networks for linear second-order elliptic and parabolic type PDEs. arXiv:2004.01806 (2020). arXiv.
- [11] Liu Yang, Xuhui Meng, George Em Karniadakis. B-PINNs: Bayesian Physics-Informed Neural Networks for Forward and Inverse PDE Problems with Noisy Data. arXiv:2003.06097 (2020). arXiv.
- [12] Balaji Lakshminarayanan, Alexander Pritzel, Charles Blundell. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. arXiv:1612.01474 (2016). arXiv.
- [13] Tongtao Zhang, Biswadip Dey, Pratik Kakkar, Arindam Dasgupta, Amit Chakraborty. Frequency-compensated PINNs for Fluid-dynamic Design Problems. ML4Eng Workshop, NeurIPS 2020. arXiv:2011.01456 (2020). arXiv.
- [14] Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, et al.. Fourier Features Let Networks Learn High Frequency Functions in Low Dimensional Domains. arXiv:2006.10739 (2020). arXiv.
- [15] Vincent Sitzmann, Julien N. P. Martel, Alexander W. Bergman, et al.. Implicit Neural Representations with Periodic Activation Functions. arXiv:2006.09661 (2020). arXiv.
- [16] Sifan Wang, Hanwen Wang, Paris Perdikaris. On the eigenvector bias of Fourier feature networks: From regression to solving multi-scale PDEs with physics-informed neural networks. arXiv:2012.10047 (2020). arXiv.
- [17] Rajat Arora, Pratik Kakkar, Amit Chakraborty, Biswadip Dey. Physics-informed neural networks for modeling rate- and temperature-dependent plasticity. Machine Learning and the Physical Sciences Workshop, NeurIPS 2022. arXiv:2201.08363v3 (2022). arXiv.
- [18] Filippo Masi, Ioannis Stefanou, Paolo Vannucci, Victor Maffi-Berthier. Thermodynamics-based Artificial Neural Networks for constitutive modeling. arXiv:2005.12183 (2020). arXiv.
- [19] Sifan Wang, Yujun Teng, Paris Perdikaris. Understanding and mitigating gradient pathologies in physics-informed neural networks. arXiv:2001.04536 (2020). arXiv.
- [20] Sifan Wang, Xinling Yu, Paris Perdikaris. When and why PINNs fail to train: A neural tangent kernel perspective. arXiv:2007.14527 (2020). arXiv.
- [21] Aditi S. Krishnapriyan, Amir Gholami, Shandian Zhe, et al.. Characterizing possible failure modes in physics-informed neural networks. arXiv:2109.01050 (2021). arXiv.
- [22] Levi McClenny, Ulisses Braga-Neto. Self-Adaptive Physics-Informed Neural Networks using a Soft Attention Mechanism. arXiv:2009.04544 (2020). arXiv.
- [23] Rafael Bischof, Michael Kraus. Multi-Objective Loss Balancing for Physics-Informed Deep Learning. arXiv:2110.09813 (2021). arXiv.
- [24] Sifan Wang, Shyam Sankaran, Paris Perdikaris. Respecting causality is all you need for training physics-informed neural networks. arXiv:2203.07404 (2022). arXiv.
- [25] Jeremy Yu, Lu Lu, Xuhui Meng, George Em Karniadakis. Gradient-enhanced physics-informed neural networks for forward and inverse PDE problems. arXiv:2111.02801 (2021). arXiv.
- [26] Chenxi Wu, Min Zhu, Qinyang Tan, et al.. A comprehensive study of non-adaptive and residual-based adaptive sampling for physics-informed neural networks. arXiv:2207.10289 (2022). arXiv.
- [27] Mohammad Amin Nabian, Rini Jasmine Gladstone, Hadi Meidani. Efficient training of physics-informed neural networks via importance sampling. arXiv:2104.12325 (2021). arXiv.
- [28] Krithika Manohar, Bingni W. Brunton, J. Nathan Kutz, Steven L. Brunton. Data-Driven Sparse Sensor Placement for Reconstruction. arXiv:1701.07569 (2017). arXiv.
- [29] Shuheng Liao, Tianju Xue, Jihoon Jeong, et al.. Hybrid thermal modeling of additive manufacturing processes using physics-informed neural networks for temperature prediction and parameter identification. arXiv:2206.07756 (2022). arXiv.

- [30] Qiming Zhu, Zeliang Liu, Jinhui Yan. Machine learning for metal additive manufacturing: Predicting temperature and melt pool fluid dynamics using physics-informed neural networks. arXiv:2008.13547 (2020). arXiv.
- [31] Tamara G. Grossmann, Urszula Julia Komorowska, Jonas Latz, Carola-Bibiane Schönlieb. Can Physics-Informed Neural Networks beat the Finite Element Method?. arXiv:2302.04107 (2023). arXiv.
- [32] Zhongkai Hao, Jiachen Yao, Chang Su, et al.. PINNacle: A Comprehensive Benchmark of Physics-Informed Neural Networks for Solving PDEs. arXiv:2306.08827 (2023). arXiv.
- [33] Makoto Takamoto, Timothy Praditia, Raphael Leiteritz, et al.. PDEBENCH: An Extensive Benchmark for Scientific Machine Learning. arXiv:2210.07182 (2022). arXiv.
- [34] Lu Lu, Xuhui Meng, Zhiping Mao, George E. Karniadakis. DeepXDE: A deep learning library for solving differential equations. arXiv:1907.04502 (2019). arXiv.
- [35] Sifan Wang, Shyam Sankaran, Hanwen Wang, Paris Perdikaris. An Expert's Guide to Training Physics-informed Neural Networks. arXiv:2308.08468 (2023). arXiv.