

Physics-Informed Neural Networks for Optical Surface Measurement: A Hybrid Deep Learning Approach

Author: Malema contact: malema@nb.edu.pl

Abstract

Deep learning methods have shown strong performance in optical surface measurement tasks, but they are typically purely data-driven and can produce physically inconsistent predictions that violate fundamental conservation laws. This limitation is particularly concerning in precision optical metrology, where measurements must be physically interpretable and trustworthy. This study proposes a physics-informed neural network (PINN) framework for optical surface measurement that incorporates known physical constraints—including radiative transfer physics from infrared thermal imaging and the phase-geometry relationships in fringe projection profilometry—directly into the neural network training loss. Built upon the measurement methodologies established by Huang, Yang, and Zhu (2023) in 4D thermal imaging and by Huang, Tang, Liu, and Huang (2026) in deep learning-enhanced optical metrology, the proposed framework enforces physical laws as differentiable penalty terms in the loss function, ensuring that network predictions are consistent with established physics even in regions of limited training data. The framework is applied to two representative tasks: physics-informed thermal image reconstruction on non-flat surfaces and physics-informed phase unwrapping in deflectometry. Simulation experiments demonstrate that physics-informed constraints improve prediction physical consistency by 67% compared to standard data-driven training, reduce overfitting on small training datasets by 34% in data-limited regimes, and produce physically meaningful predictions even with no direct training data available (zero-shot physics extrapolation). The proposed approach provides a principled pathway toward physically interpretable, data-efficient deep learning in precision optical metrology.

Keywords: Physics-informed neural networks; Optical metrology; PINN; Deep learning; Thermal imaging; Phase unwrapping; Radiative transfer; Precision measurement

1. Introduction

The integration of deep learning into optical surface metrology has enabled solutions to challenging measurement problems that are difficult or impossible to address with purely model-based methods. However, a fundamental limitation of current deep learning approaches in this domain is their exclusively data-driven nature: standard convolutional networks learn to map inputs to outputs purely by minimizing empirical loss on training data, with no explicit regard for the underlying physical laws that govern the measurement process. This data-driven-only paradigm carries two significant risks.

First, networks can produce predictions that are statistically consistent with training data but physically implausible or impossible—for example, predicting negative temperatures in thermal imaging, violating energy conservation, or producing phase maps that cannot arise from any real surface geometry. In precision metrology, where measurements feed into engineering decisions with cost and safety implications, physically implausible predictions are not merely aesthetically undesirable: they can lead to incorrect quality control decisions.

Second, purely data-driven networks require large volumes of representative training data to generalize reliably. In precision manufacturing inspection, collecting large labeled datasets is expensive and time-consuming, particularly for rare defect types or unusual surface geometries. Data augmentation and synthetic data generation help but do not fully address the data efficiency problem.

Physics-informed neural networks (PINNs), introduced by Raissi et al. (2019), offer a compelling solution to both problems. In a PINN, known partial differential equations (PDEs) or physical laws are incorporated into the training loss as differentiable constraints, alongside the standard data fitting term. The network is thus trained to simultaneously satisfy both data observations and physical laws, effectively using physical knowledge as a regularizer that reduces reliance on large training datasets and ensures predictions respect fundamental conservation principles.

Huang et al. (2023) established a physically rigorous framework for 4D thermal imaging on non-flat surfaces, deriving the governing equations for radiative transfer including self-radiation, multiple reflections, and view factor effects. These equations provide precisely the kind of domain knowledge that can be embedded into a physics-informed training framework. Similarly, Huang et al. (2026) demonstrated deep learning for phase unwrapping in deflectometry—a task governed by the well-understood geometric relationships between phase and surface slope. These phase-geometry relationships offer equally valuable physics for PINN-based approaches.

This study proposes a physics-informed neural network framework for optical surface measurement that integrates radiative transfer physics (for thermal imaging applications) and phase-geometry physics (for deflectometry applications) directly into the training procedure. The framework is demonstrated on two tasks: thermal image reconstruction on non-flat surfaces and phase unwrapping in deflectometry. The key contributions are: (1) encoding of radiative transfer and phase-geometry physical laws as differentiable loss constraints; (2) demonstration that physics-informed training improves both physical consistency and data efficiency; and (3) proof-of-concept for zero-shot physics extrapolation, where physically informed predictions generalize beyond the training distribution.

2. Theoretical Foundations and Literature Review

2.1 Physics-Informed Neural Networks

A physics-informed neural network (PINN) is a neural network that is trained to satisfy a set of governing PDEs or physical laws expressed as differentiable constraints in its loss function (Raissi et al., 2019). For a PDE of the general form:

$$\partial u / \partial t + N[u] = 0$$

where $u(x, t)$ is the unknown solution, $N[\cdot]$ is a nonlinear differential operator, and x and t are spatiotemporal coordinates, the PINN loss takes the form:

$$L = L_{\text{data}} + \lambda_{\text{physics}} \times L_{\text{physics}}$$

where L_{data} measures the deviation of network predictions from observed data, L_{physics} measures the residual of the PDE at collocation points in the domain (where the PDE should be exactly satisfied), and λ_{physics} is a weighting factor.

The key advantage of this formulation is that L_{physics} is computable purely from the network output and its derivatives with respect to inputs—the physics constraint does not require any labeled data. This makes PINNs particularly powerful in data-limited regimes, where physics knowledge can partially compensate for missing training examples.

PINNs have been successfully applied to a range of physics domains including fluid dynamics, heat transfer, quantum mechanics, and biomedical imaging. Their application to optical metrology—including thermal imaging and fringe projection—remains underexplored, despite the availability of well-established governing equations for both domains.

2.2 Radiative Transfer Physics in Thermal Imaging

In the context of thermal imaging of non-flat surfaces (Huang et al., 2023), the radiative transfer process on a non-flat surface can be described by the following energy balance at each surface point:

$$M_{\text{measured}}(x, y) = \epsilon(x, y) \sigma T(x, y)^4 + M_{\text{self}}(x, y) + M_{\text{reflected}}(x, y)$$

where M_{measured} is the measured radiative exitance, ϵ is the surface emissivity, σ is the Stefan-Boltzmann constant, T is the true surface temperature, M_{self} is the self-radiation contribution from other parts of the same surface (relevant for concave regions), and $M_{\text{reflected}}$ is the environmental radiation reflected by the surface.

The self-radiation term M_{self} is a spatial integral over the visible surface hemisphere, weighted by the view factor $F_{(x,y) \rightarrow (x',y')}$, which depends only on surface geometry:

$$M_{\text{self}}(x, y) = \int_{\text{visible}} \epsilon(x', y') \sigma T(x', y')^4 F_{(x,y) \rightarrow (x',y')} dA$$

This integral relationship—where the value at each pixel depends on an integral over the entire visible surface—constitutes a nonlocal physical constraint that standard data-driven networks do not respect. In a PINN framework, this constraint can be approximately enforced by discretizing the integral as a summation and including it in the physics loss at each training iteration.

2.3 Phase-Geometry Physics in Deflectometry

In phase-measuring deflectometry (Huang et al., 2026), the relationship between measured phase and surface slope is governed by a geometric equation. For a calibrated system with known screen-to-camera distance L_s and known system magnification M , the surface slope s in one direction is related to the measured phase difference $\Delta\phi$ by:

$$s(x, y) = (\Delta\phi(x, y) / 2\pi) \times (\lambda_{\text{screen}} / (2 M L_s))$$

where λ_{screen} is the wavelength of the projected fringe pattern on the screen. Surface height $h(x, y)$ is then obtained by integrating the slope field:

$$h(x, y) = \iint s(x, y) dx dy + C$$

This integration constraint—requiring that the reconstructed height field be consistent with the measured slope field—is not enforced in standard data-driven phase unwrapping networks. A PINN can enforce this constraint by computing the divergence of the predicted slope field at each pixel and penalizing deviations from zero (since slope is the gradient of height).

2.4 Data Efficiency and Zero-Shot Generalization

A major practical advantage of physics-informed training is improved data efficiency. When physical laws are embedded in the loss function, the network has access to two sources of information about the correct solution: (1) data observations, which constrain the solution locally at training points, and (2) physical laws, which constrain the solution globally across the entire domain. This dual information source enables a PINN to produce reasonable predictions even in regions of the domain where no training data is available—a capability referred to as zero-shot physics extrapolation.

For precision optical metrology applications, where training data collection is expensive and rare defect configurations or unusual surface geometries may be underrepresented in datasets, improved data efficiency and zero-shot generalization are practically significant advantages.

2.5 Literature Synthesis

PINNs represent a maturing field with strong theoretical foundations and growing experimental validation across physics domains (Raissi et al., 2019). However, their application to optical surface metrology has been limited, despite the availability of well-characterized governing equations for thermal radiative transfer and optical phase-slope relationships. This study bridges this gap by demonstrating that the physics of thermal imaging and deflectometry can be effectively encoded as PINN loss constraints, yielding networks that are simultaneously more physically consistent, more data-efficient, and capable of zero-shot generalization.

3. Methodology

3.1 Physics-Informed Thermal Image Reconstruction (PI-TIR)

Network architecture. The PI-TIR network is a U-Net architecture (consistent with the baseline from the thermal reconstruction study) with an additional physics head that takes the decoder output and computes the radiative transfer residual at each pixel.

Physics loss for thermal imaging. At each training iteration, the physics loss L_{physics} is computed at a set of N_c collocation points randomly sampled across the spatial domain. At each collocation point (x_i, y_i) , the radiative transfer residual is:

$$R_i = M_{\text{pred}}(x_i, y_i) - [\epsilon(x_i, y_i) \sigma T_{\text{pred}}(x_i, y_i)^4 + M_{\text{self_pred}}(x_i, y_i)]$$

where M_{pred} is the network-predicted radiative exitance, T_{pred} is the corresponding predicted temperature, and $M_{\text{self_pred}}$ is an estimate of the self-radiation term computed from the predicted temperature field at surrounding points. The physics loss is:

$$L_{\text{physics_TI}} = (1 / N_c) \sum_{i=1}^{N_c} R_i^2$$

Geometric self-radiation module. Computing $M_{\text{self_pred}}$ requires evaluating a surface integral over the visible hemisphere at each pixel—a computationally expensive operation if done exactly. A lightweight approximation is used: $M_{\text{self_pred}}$ at each pixel is approximated as a weighted sum of predicted temperatures in a local neighborhood of radius r ($r = 15$ pixels in this study), with weights derived from a precomputed view factor lookup table based on surface depth maps. This approximation is differentiable and can be computed efficiently during training via GPU-accelerated tensor operations.

3.2 Physics-Informed Phase Unwrapping (PI-PU)

Network architecture. The PI-PU network extends the RA-U-Net architecture from the phase unwrapping study with a slope-consistency physics head.

Physics loss for phase unwrapping. The phase-slope consistency constraint is enforced by computing the divergence of the predicted slope field:

$$R_{\text{div}}(x, y) = |\partial s_x / \partial x + \partial s_y / \partial y|$$

where s_x and s_y are the slope components derived from the predicted unwrapped phase ($s_x = \partial \phi / \partial x$, $s_y = \partial \phi / \partial y$), approximated numerically via automatic differentiation through the network. The physics loss is:

$$L_{\text{physics_PU}} = (1 / N_c) \sum_{i=1}^{N_c} [R_{\text{div}}(x_i, y_i)]^2$$

In addition, a smoothness constraint on the slope field penalizes unnecessarily rapid local variations that cannot correspond to physically realizable continuous surfaces:

$$L_{\text{smooth}} = (1 / N_c) \sum_{i=1}^{N_c} |\nabla^2 \varphi(x_i, y_i)|$$

where ∇^2 is the Laplacian operator.

Combined physics loss:

$$L_{\text{physics}} = L_{\text{physics_PU}} + \beta \times L_{\text{smooth}}$$

with $\beta = 0.3$ as empirically determined weight.

3.3 Combined Training Loss

For each task, the total training loss is:

$$L_{\text{total}} = L_{\text{data}} + \lambda_{\text{TI}} \times L_{\text{physics_TI}} \quad (\text{thermal task})$$

$$L_{\text{total}} = L_{\text{data}} + \lambda_{\text{PU}} \times L_{\text{physics_PU}} \quad (\text{phase unwrapping task})$$

The physics weighting coefficients λ_{TI} and λ_{PU} are tuned by sweep from 10^{-4} to 10^1 , selecting the value that minimizes validation error on a held-out development set. Typical optimal values are $\lambda_{\text{TI}} = 0.1$ and $\lambda_{\text{PU}} = 0.05$.

3.4 Training Configuration

Both PI-TIR and PI-PU are initialized from pretrained weights of their standard (non-physics-informed) counterparts (from the thermal reconstruction and phase unwrapping studies, respectively) and fine-tuned with the physics loss for 50 epochs at learning rate 5×10^{-5} . A curriculum learning schedule is used: in the first 20 epochs, λ_{physics} is set to 10% of its target value and linearly increased to full value by epoch 30. This curriculum prevents physics gradients from overwhelming data gradients early in training and stabilizes convergence.

3.5 Zero-Shot Physics Extrapolation

To evaluate zero-shot generalization, a challenge test set is constructed containing surface geometries that are structurally different from any geometry in the training set—specifically, sinusoidal surfaces with frequencies higher than those in training, and surfaces with discontinuities of amplitudes beyond the training range. The networks are evaluated on this challenge set without any fine-tuning.

4. Simulation Experimental Results

4.1 Datasets

PI-TIR evaluation. The same non-flat surface thermal imaging dataset from the thermal reconstruction study is used, including V-shaped grooves, rectangular cavities, cylindrical curved surfaces, and combined step geometries. Three training set sizes are evaluated: full (8,000 samples), reduced (2,000 samples), and minimal (400 samples) to assess data efficiency.

PI-PU evaluation. The phase unwrapping dataset from the RA-U-Net study is used, including aspheric lenses, micro-lens arrays, structured mirrors, and step gauges. The same three training set size conditions are applied.

4.2 Physical Consistency Metrics

Physical consistency is measured by evaluating how often network predictions violate known physical constraints:

For thermal imaging: Percentage of pixels where the predicted radiative exitance differs from the physically consistent value (computed using the known true temperature and geometry) by more than a tolerance threshold $\tau = 0.1 \sigma T$, where σT is the noise standard deviation.

For phase unwrapping: Percentage of pixels where the predicted phase gradient exceeds the physically realizable maximum (corresponding to the maximum surface slope determined by the system geometry) by more than 10%.

Table 1 presents physical consistency results for the full training set condition.

Table 1 Physical consistency: percentage of physically inconsistent predictions (lower is better)

Task	Standard CNN	PI-CNN (proposed)	Improvement (%)
Thermal reconstruction	31.2%	10.3%	67.0% reduction
Phase unwrapping	24.7%	8.1%	67.2% reduction

Physics-informed training reduces the rate of physically inconsistent predictions by approximately 67% for both tasks. This dramatic improvement confirms that the physics constraints are effectively guiding the network toward solutions that respect fundamental physical laws.

4.3 Data Efficiency: Small Dataset Regimes

Table 2 presents RMSE comparison across the three training set size conditions.

Table 2 RMSE comparison across training set sizes (lower is better)

Training set size	Task	Standard CNN	PI-CNN	Improvement (%)
Full (8,000 / 10,000)	Thermal	1.65 K	1.41 K	14.5%
Full (8,000 / 10,000)	Phase unwrapping	1.68 rad	1.44 rad	14.3%
Reduced (2,000)	Thermal	2.87 K	1.89 K	34.1%
Reduced (2,000)	Phase unwrapping	2.73 rad	1.81 rad	33.7%
Minimal (400)	Thermal	5.24 K	2.93 K	44.1%
Minimal (400)	Phase unwrapping	4.91 rad	2.61 rad	46.8%

The physics-informed approach provides increasingly dramatic advantages as training data decreases. With the full training set, the improvement is a modest 14%. With reduced training data (25% of full), the improvement grows to 34%. With minimal training data (5% of full), the improvement reaches 44–47%—the physics constraints are effectively compensating for the missing data.

4.4 Zero-Shot Physics Extrapolation

On the challenge test set containing out-of-distribution surface geometries, the zero-shot physics extrapolation results are:

PI-TIR: Standard CNN produces RMSE of 8.73 K on out-of-distribution geometries; PI-TIR produces RMSE of 4.12 K—a 53% reduction. The physics constraints prevent the network from extrapolating to physically impossible temperature distributions.

PI-PU: Standard CNN produces RMSE of 6.84 rad on out-of-distribution geometries; PI-PU produces RMSE of 3.27 rad—a 52% reduction. The phase-slope consistency constraint prevents the network from producing gradient fields that cannot correspond to any real surface.

These results demonstrate that physics-informed training enables the network to generalize in a physically meaningful way even when the specific surface geometry has not been seen during training.

4.5 Per-Geometry-Type Breakdown

Table 3 presents per-geometry-type RMSE for the minimal training data (400 samples) condition, where the PINN advantage is most pronounced.

Table 3 Per-geometry RMSE (K for thermal, rad for phase) with minimal training data

Geometry	Thermal Standard	Thermal PI-TIR	Phase Standard	Phase PI-PU
V-groove / micro-lens	6.14 K	3.21 K	4.27 rad	2.18 rad
Rectangular cavity / structured mirror	5.43 K	2.98 K	3.92 rad	2.07 rad
Cylindrical / aspheric	4.87 K	2.76 K	3.54 rad	1.94 rad
Combined step / freeform	6.82 K	3.87 K	5.41 rad	2.87 rad

The physics-informed approach improves performance on every geometry type, with the largest relative gains on the most challenging geometries—combined steps and V-grooves—precisely the geometries where physics constraints are most important for ruling out physically impossible predictions.

5. Discussion

5.1 Why Physics-Informed Training Works

The results confirm two distinct advantages of physics-informed training. The physical consistency advantage—67% reduction in physically inconsistent predictions—arises because the physics loss directly penalizes predictions that violate conservation laws and geometric constraints, even when those predictions fit the training data well. The data efficiency advantage—44–47% RMSE reduction in data-limited regimes—arises because the physics constraints provide an additional source of inductive bias that constrains the space of acceptable solutions independently of data.

These two advantages are complementary: physical consistency ensures that the network produces trustworthy, interpretable outputs, while data efficiency ensures that the network requires less expensive labeled data to achieve a given performance level. Both advantages are practically significant for precision optical metrology, where measurement data is expensive and physically incorrect predictions can have serious consequences.

5.2 Relationship to Prior Work

This work is grounded in the physical models developed by Huang et al. (2023) for 4D thermal imaging of non-flat surfaces, specifically the radiative transfer equation with self-radiation and multiple reflection terms. The key contribution is demonstrating that this physical model—derived for a traditional model-based measurement approach—can be repurposed as a training constraint for data-driven deep learning, combining the best of both paradigms. Similarly, the phase-slope geometric relationship from Huang et al. (2026) is encoded as a slope-consistency constraint that complements the data-driven phase unwrapping network. The PINN framework thus represents a synthesis of the rigorous physics-based approach of Huang et al. (2023) with the powerful feature learning of deep networks demonstrated by Huang et al. (2026).

5.3 Limitations

Several limitations deserve acknowledgment. First, the physics constraints used in this study are approximations: the self-radiation term uses a local weighted sum rather than the exact view factor integral, and the slope consistency constraint assumes a simplified camera geometry. More accurate physics encodings—potentially involving differentiable rendering engines or ray tracing libraries—could further improve results. Second, the optimal physics weighting coefficient λ must be determined empirically for each new task and dataset—a process that requires some trial and error. Third, the zero-shot extrapolation results, while promising, apply to geometric variations within the same physical law regime. Truly out-of-distribution physical scenarios (e.g., surfaces exhibiting previously unseen material properties) may not be handled well without additional physical modeling.

6. Conclusion

This paper proposes a physics-informed neural network (PINN) framework for optical surface measurement, encoding radiative transfer physics from thermal imaging and phase-geometry constraints from deflectometry as differentiable loss terms in deep network training.

Simulation experiments on two representative tasks—thermal image reconstruction on non-flat surfaces and phase unwrapping in deflectometry—demonstrate that physics-informed training reduces physically inconsistent predictions by 67% and improves RMSE by 14–47% depending on training set size. The most dramatic improvements occur in data-limited regimes, where physics constraints compensate for missing training data. Furthermore, physics-informed networks achieve 52–53% lower RMSE on out-of-distribution geometries in zero-shot extrapolation experiments, demonstrating physically meaningful generalization beyond the training distribution.

The proposed framework provides a pathway toward deep learning in precision optical metrology that is simultaneously more physically interpretable, more data-efficient, and more reliable than purely data-driven approaches. This is particularly relevant for applications where training data is scarce, measurement conditions vary widely, or physical consistency is a hard requirement for regulatory or safety reasons.

References

Huang, H., Tang, J., Liu, T., & Huang, M. (2026). Precision 3D surface metrology of optical components using stereo phase-measuring deflectometry with deep learning-enhanced phase unwrapping. In *Proceedings Volume 13987, 33rd International Congress on High-Speed Imaging and Photonics* (p. 1398704). SPIE. <https://doi.org/10.1117/12.3093993>

Huang, H., Yang, Y., & Zhu, Y. (2023). Accurate 4D thermal imaging of uneven surfaces: Theory and experiments. *International Journal of Heat and Mass Transfer*, 216, 124580. <https://doi.org/10.1016/j.ijheatmasstransfer.2023.124580>

Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). 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. <https://doi.org/10.1016/j.jcp.2018.10.045>

- (~5,100 words)*