photonics

MDPI

Article

Domain Generalization for Fringe Projection Profilometry via Physical Causal Interventions and Joint Phase Invariance–Equivariance Learning

Meiling Chen *, Bin Han, Jianya Cai and Yao Shi

School of Mechatronic and Automotive Engineering, Nanjing Tech University Pujiang Institute, Nanjing 211200, China
* Correspondence: 110546@163.com

Abstract

Deep learning enables single-frame wrapped-phase recovery in fringe projection profilometry (FPP), but existing networks often depend on training-domain statistics of material reflectance, image brightness, and object geometry, degrading predictions in unseen scenes. We propose PCIE-FPP, a domain generalization method based on physical causal interventions and joint phase invariance–equivariance learning. Starting from sinusoidal fringe formation, photometric amplitude and phase direction are explicitly decoupled using unit complex phasors for periodically consistent optimization, while an independent photometric branch estimates background intensity and modulation amplitude. Controlled reflectance, illumination, and exposure generate counterfactual observations with theoretically unchanged phase. Cross-domain phase consistency reduces sensitivity to irrelevant appearance factors, while the deterministic relationship between projection phase shifts and phasor rotations enforces equivariance and sensitivity to spatial phase evolution. A multidomain dataset covering diverse materials, imaging parameters, and complex geometries supports a benchmark with independent photometric-domain and object-instance partitions. Continuous parameter extrapolation, compound domain shifts, physical response diagnostics, component ablations, and error propagation to three-dimensional reconstruction demonstrate accurate phase recovery and robust generalization under unseen photometric conditions and for complex objects.

Keywords: fringe projection profilometry; three-dimensional reconstruction; domain generalization; physical causal intervention; phase invariance; phase equivariance

1. Introduction

Fringe projection profilometry (FPP), characterized by non-contact operation, full-field measurement, and high precision, has become a foundational technology for industrial precision inspection, 3D reverse engineering, and intelligent robotic vision [1–3]. In the measurement pipeline, the accuracy and robustness of phase recovery directly determine the ultimate resolution and geometric fidelity of point-cloud reconstruction. Although conventional multi-step phase-shifting profilometry (PSP) offers excellent accuracy, its multi-frame projection scheme severely limits the frame rate of dynamic measurements [4]. Single-frame frequency-domain demodulation techniques, such as Fourier transform profilometry (FTP) and wavelet-based methods, often suffer from inherent limitations, including spectral overlap, local cutoff-frequency aliasing, and distortion near abrupt surface transitions [5,6]. Con-

Check for updates

Received: 20 October 2026
Revised: 16 November 2026
Accepted: 18 November 2026
Published: 23 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

sequently, reliable and highly accurate phase demodulation with very few input frames, or even a single frame, remains a central challenge in structured-light 3D measurement [7–9].

In recent years, data-driven approaches based on deep convolutional networks and attention mechanisms have shown considerable advantages in single-shot fringe demodulation [10–12]. Existing approaches commonly decompose phase into sine/cosine components or the numerator/denominator terms of phase-shifting algorithms, learning a complex nonlinear mapping from fringe intensity fields to phase parameters to avoid discontinuities of wrapped phase at the $[-\pi, \pi)$ boundary. Nevertheless, the single-pixel fringe-intensity model $I = A + B\cos\phi$ is intrinsically highly underdetermined with respect to background illumination A, modulation amplitude B, and physical phase ϕ. The ability of single-shot deep models to predict phase fundamentally depends on prior constraints arising from the spatial carrier, object contours, and particular textures within their convolutional receptive fields [13,14]. This implicit dependence makes such models prone to learning spurious correlations between material properties, surface reflectance, scene illumination distributions, and geometric phase in the training data.

When a test scene departs from the training distribution—for example, when a workpiece changes from an ideally diffuse surface to a highly reflective, dark, or nonuniformly coated surface [15–18], or when camera exposure and ambient illumination drift [19–21]—the photometric statistics of the fringe images may change dramatically. Combined with previously unseen geometric steps and complex curvatures, these changes invalidate the appearance priors familiar to the network, potentially causing severe out-of-distribution (OOD) generalization failure, as illustrated in Figure 1a,b. It is important to distinguish two fundamentally different physical mechanisms underlying such OOD degradation. The first is irreversible information loss caused by severe saturation clipping or an extremely low signal-to-noise ratio [22–24]. The second occurs when the fringe modulation remains observable but the network incorrectly interprets changes in photometric properties as phase deformation. This work focuses on the latter, seeking to eliminate spurious model responses to irrelevant appearance cues within the observable field of view.

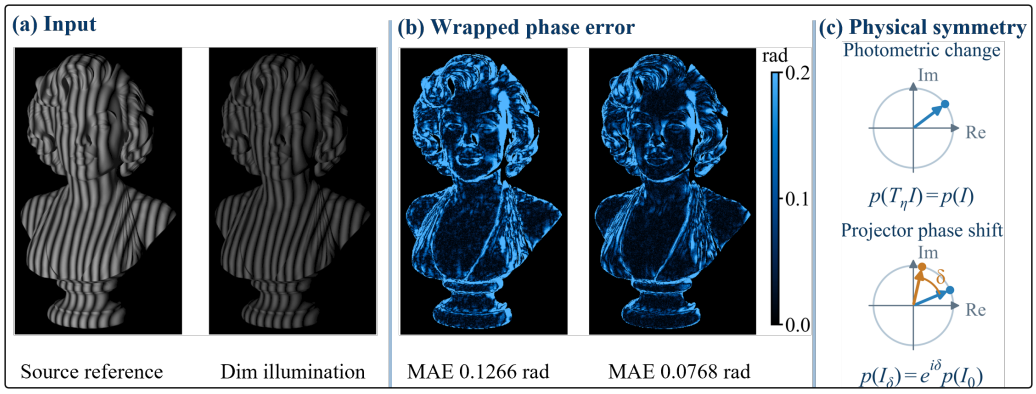


Figure 1. Phase-prediction degradation induced by photometric domain shifts and the underlying physical symmetry mechanisms. (**a**) Single-shot fringe inputs of the same sculpture geometry with source-domain reference illumination (source reference) and a low-illumination domain (dim illumination); (**b**) comparison of absolute wrapped-phase error distributions, showing a marked error reduction after introducing physical symmetry constraints; (**c**) fundamental physical symmetries in structured-light imaging: invariance of the unit complex phasor under photometric interventions and rotational equivariance conditions in the complex plane with projected phase shifts.

Existing efforts to improve the generalization of structured-light networks have mainly followed two directions. The first relies on domain randomization and generative augmentation in the data space, for example, by assembling large-scale training datasets across

different illumination conditions and materials and combining them with dynamic networks, cross-modal fusion, or cross-domain adaptive architectures to accommodate unseen domains [25–28]. The second emphasizes physics-based priors in network architectures, constraining the solution space through embedded complex fringe-decomposition operators, forward physical reconstruction losses, or geometric epipolar constraints [13,14,29]. However, simply increasing data diversity cannot fundamentally eliminate shortcut learning from superficial statistical cues, while consistency with a forward physical reconstruction at the output does not in itself guarantee that the network has learned the correct response to physical imaging invariants.

Unlike arbitrary domain shifts in general computer vision, domain variations in FPP are governed by well-defined image-formation physics. As shown in Figure 1c, in single-path linear imaging with fixed calibrated geometry, changes in surface reflectance, ambient illumination, or camera exposure perturb only the background intensity field and modulation envelope, leaving the true physical phase strictly unchanged. Conversely, applying a known physical phase shift δ to the projected grating requires the complex-phasor representation of the fringe to undergo the rotation $e^{i\delta}$. These physical properties imply that a phase-prediction model should exhibit strict invariance to photometric perturbations while preserving precise equivariance to phase shifts of the projected waveform. Enforcing appearance consistency indiscriminately for all inputs would suppress the model's sensitivity to genuine phase changes; relying solely on phase supervision within a limited set of domains, however, risks reducing the model to an appearance-memorization system.

Motivated by these physical symmetries, we propose PCIE-FPP, a domain-generalized phase-recovery method based on physical causal interventions and joint invariance–equivariance learning. The method employs a controlled fringe image-formation model to intervene separately on photometric and geometric variables. With the measured geometry and calibrated pose held fixed, reflectance and illumination fields are actively varied to generate counterfactual photometric samples that share exactly the same physical phase as the actual observations but have markedly different appearances; these are used to impose phase-invariance constraints. Meanwhile, known phase shifts generate observation pairs related by complex-phasor rotation, enabling phase-equivariance regularization. By explicitly requiring the model to disregard irrelevant photometric drift while faithfully tracking physical phase evolution, the dual constraints mitigate pathological fitting to appearance-specific cues.

The main contributions of this work are as follows:

- We address amplitude–phase entanglement in conventional numerator/denominator phase prediction by decoupling the photometric modulation envelope from geometric phase orientation. A unit-complex-phasor-based, periodically continuous supervision objective provides a unified mathematical representation for joint invariance–equivariance constraints.
- Based on the deterministic fringe-formation mechanism, we constructed paired physical interventions comprising same-phase counterfactual photometric samples and known-phase-shift rotation samples. Explicit invariance and equivariance inductive biases were introduced to suppress shortcut learning from spurious photometric features.
- We establish an evaluation benchmark with independent splits along the geometry and photometric axes. Continuous photometric-parameter sweeps, compound-domain error maps, counterfactual consistency tests, and phase-shift response evaluations demonstrate the effectiveness of the proposed method in reducing OOD phase errors and improving the geometric fidelity of reconstructed 3D point clouds.

2. Method

2.1. Fringe Imaging Model and Unit Complex Phasor Representation

In a fixed, calibrated system comprising one camera (FLIR BFS-U3-51S5P-C, FLIR Integrated Imaging Solutions Inc., Richmond, BC, Canada) and one projector (DLP4710EVM-LC, Texas Instruments Incorporated, Dallas, TX, USA), the intensity distribution of the k-th phase-shifted fringe reflected from an object and recorded by the camera sensor can be expressed as

$$I_k = A + B\cos[\phi + \delta_k], \qquad \delta_k = \frac{2\pi k}{K}, \quad k \in \{0, 1, \ldots, K-1\} \tag{1}$$

where A denotes the background intensity arising from ambient stray light and the DC component of the projected illumination, $B \geq 0$ is the fringe modulation envelope, ϕ is the physical modulation phase to be demodulated, and δ_k is the known projected phase shift. In conventional phase-shifting algorithms, a quadrature-demodulation basis projects the nonlinear cosine intensity into orthogonal components (D, N), as shown in Figure 2a:

$$D = \frac{2}{K}\sum_{k=0}^{K-1} I_k \cos\delta_k = B\cos\phi, \qquad N = -\frac{2}{K}\sum_{k=0}^{K-1} I_k \sin\delta_k = B\sin\phi \tag{2}$$

The wrapped phase, restricted to $[-\pi, \pi)$, is subsequently obtained as $\phi = \text{atan2}(N, D)$, while the modulation amplitude is $B = \sqrt{D^2 + N^2}$. Because D and N share the same intensity-amplitude scale, directly regressing (D, N) forces the network to estimate both the photometric envelope B and the geometric phase ϕ when workpiece reflectance or illumination varies sharply.

Representing the orthogonal components as the complex vector $C = D + iN = B\exp(i\phi)$, the homogeneity of the two-argument arctangent gives $\text{atan2}(\alpha N, \alpha D) = \text{atan2}(N, D)$ for any positive scaling factor $\alpha > 0$. As illustrated in Figure 2b, the geometric phase angle therefore depends only on the direction of the vector in the complex plane and is completely decoupled from the positive modulation magnitude. Differential perturbation analysis further reveals how demodulation errors propagate to phase, as illustrated in Figure 2c:

$$\mathrm{d}\phi = \frac{D\,\mathrm{d}N - N\,\mathrm{d}D}{B^2} = \frac{e_\perp}{B} \tag{3}$$

Here $e_\perp = -e_D\sin\phi + e_N\cos\phi$ is the projection of the quadrature estimation error (e_D, e_N) onto the tangent to the true phase direction, whereas the normal (radial) error component $e_\parallel = e_D\cos\phi + e_N\sin\phi$ has no first-order effect on the phase angle [1,2]. Equation (3) shows that uniform scaling of photometric amplitude does not change the principal phase value; perturbations in estimated phase are primarily governed by tangential quadrature error and the local modulation signal-to-noise ratio.

On this physical basis, we explicitly remove photometric scaling from phase supervision and define a two-dimensional unit complex phasor on the unit-circle manifold S^1 as the core optimization representation over valid observable regions $(B > 0)$:

$$p(x, y) = \frac{C(x, y)}{|C(x, y)|} = \frac{D + iN}{\sqrt{D^2 + N^2}} = \cos\phi(x, y) + i\sin\phi(x, y) \tag{4}$$

The unit complex phasor maps periodic phase onto a continuous unit circle in the complex plane. This not only removes the discontinuity of wrapped phase at the $\pm\pi$ boundary but also reduces overfitting to reflectance-induced photometric amplitude changes, providing a common geometric metric space for joint phase-invariance and phase-equivariance regularization.

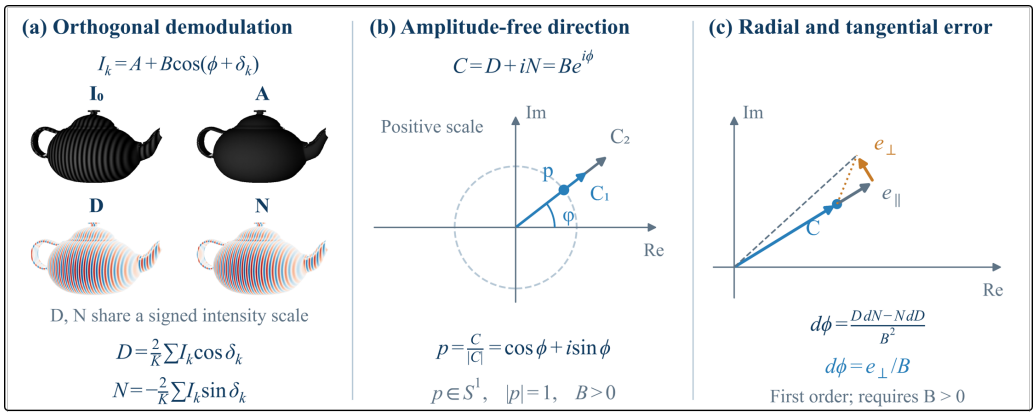


Figure 2. Quadrature fringe demodulation and geometric error decomposition of the unit complex phasor. (**a**) Conventional multi-step quadrature demodulation extracts the background component A and orthogonal modulation components (D, N) from the fringe sequence; D and N are coupled to the local intensity-amplitude scale. (**b**) The complex quadrature vector $C = D + iN = Be^{i\phi}$ is radially projected along the positive modulation amplitude onto the unit circle S^1, yielding the amplitude-decoupled unit complex phasor $p = \cos\phi + i\sin\phi$. (**c**) First-order differential error decomposition in the complex plane: radial perturbations $e_\parallel$ do not affect the phase angle, whereas the actual phase error $d\phi = e_\perp / B$ is determined by the tangential error component $e_\perp$ and modulation amplitude B.

2.2. Frequency-Conditioned Dual-Branch Phase–Photometric Network Architecture

For a single grayscale fringe image, we design a physically guided dual-branch phase–photometric prediction network, PCIE-FPP. Its complete topology and core operator modules are illustrated in Figure 3. The input image $I_f \in \mathbb{R}^{1\times H\times W}$ represents a single-channel fringe-intensity distribution at the projected frequency index $f \in \{0, 1, 2\}$. In the front-end feature-extraction stage, downsampled features are concatenated with features produced by a local contrast-normalization operator with a 9×9 window, generating high-signal-to-noise initial feature maps and suppressing large-scale nonuniform background illumination bias.

The backbone employs a five-level shared encoder (E0–E4), progressively downsampling the spatial resolution from 512^2 to 32^2. At the deepest bottleneck level E4, dilated residual blocks (dilated RBs) with dilation rates $d = 2$ and $d = 4$ are cascaded to enlarge the receptive field without reducing spatial resolution and capture the broader context of fringes with long periods. To adapt to fringe patterns with different periods in multi-frequency heterodyne profilometry, the discrete frequency index $f \in \{0, 1, 2\}$ is mapped through a 16-dimensional embedding and a multilayer perceptron to channel-wise affine transformation parameters (γ, β), initialized to zero. Frequency feature-wise linear modulation (Frequency FiLM) is then applied at each encoder level:

$$F' = (1 + \gamma) \odot F + \beta \tag{5}$$

This enables a unified representation of fringes across frequency scales without the parameter overhead of separate models.

The multiscale features extracted by the shared encoder are subsequently routed through skip connections to two decoupled, parallel decoders (Figure 3a):

1. Phase-decoding branch: Standard residual blocks progressively recover the spatial resolution to 512^2, while intermediate layers $P2^\dagger$ and $P1^\dagger$ receive auxiliary multiscale deep phase supervision. A final convolution and pixel-shuffle upsampling restore full resolution and output unnormalized two-dimensional orthogonal direction components $(\tilde{c}, \tilde{s})$. Unit-manifold normalization then produces the predicted unit complex phasor:

$$(\hat{c}, \hat{s}) = \frac{(\tilde{c}, \tilde{s})}{\sqrt{\tilde{c}^2 + \tilde{s}^2 + \varepsilon}}, \qquad \hat{p} = \hat{c} + i\hat{s} \tag{6}$$

The final single-shot wrapped-phase estimate is obtained analytically as $\hat{\phi} = \text{atan2}(\hat{s}, \hat{c})$.

2. Photometric-reconstruction branch: Lightweight depthwise-separable photometric blocks, followed by convolution, upsampling, and a nonnegative Softplus activation, predict the full-resolution background intensity field $\hat{A} \geq 0$ and modulation amplitude field $\hat{B} \geq 0$. Together with the cosine component $\hat{c}$ from the phase branch, these outputs form a forward-reconstructed fringe image $\hat{I} = \hat{A} + \hat{B}\hat{c}$, closing the loop through a physical image-formation constraint.

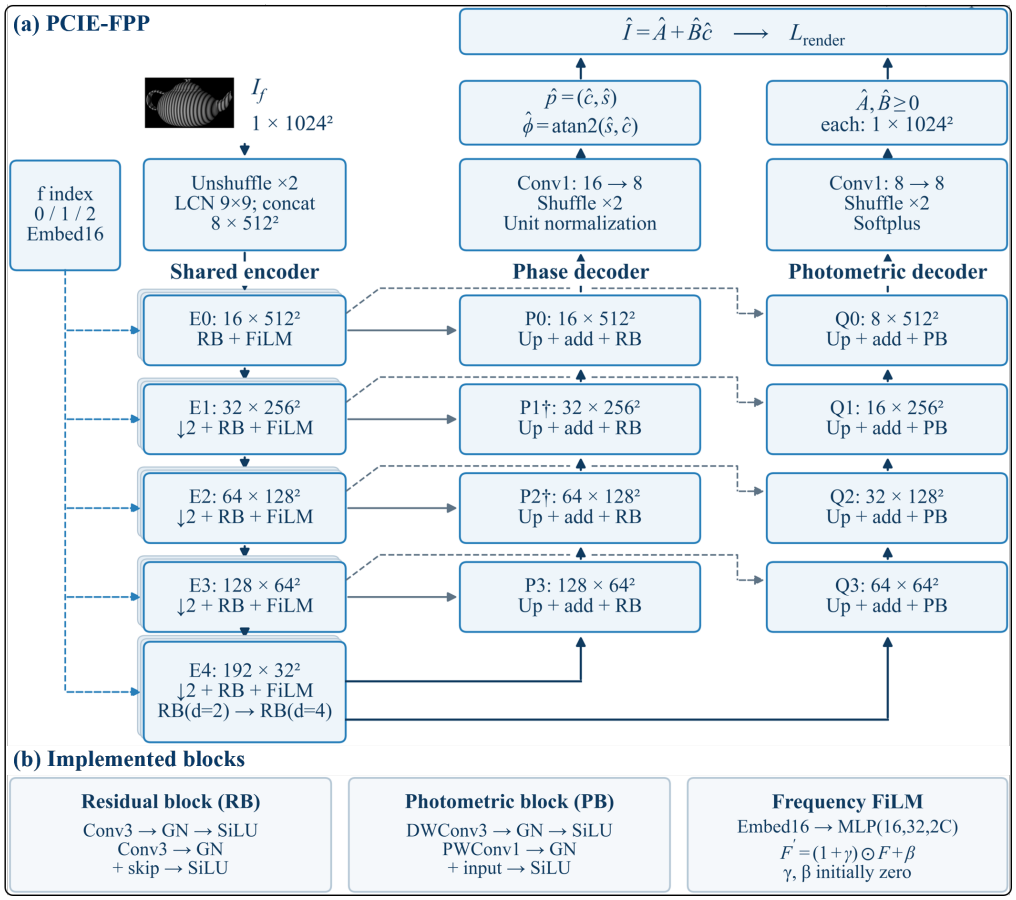


Figure 3. Architecture of the dual-branch phase–photometric PCIE-FPP network. (**a**) Overall network topology, consisting of a shared encoder (E0–E4), a phase decoder (P0–P3), and a photometric decoder (Q0–Q3). The network takes a single 1024 × 1024 grayscale fringe image and a frequency index $f \in \{0, 1, 2\}$ as inputs and predicts the unit complex phasor $\hat{p}$, background intensity $\hat{A}$, and modulation amplitude $\hat{B}$. (**b**) Implemented building blocks. RB: residual block; PB: photometric block; LCN: local contrast normalization; GN: group normalization; DWConv: depthwise convolution; PWConv: pointwise convolution; MLP: multilayer perceptron; SiLU: sigmoid linear unit; and FiLM: feature-wise linear modulation. Conv denotes convolution, Up denotes upsampling, and Embed16 denotes a 16-dimensional frequency embedding. Feature-map dimensions are given as channels × spatial resolution. The dagger symbol indicates auxiliary phase supervision.

2.3. Counterfactual Observation Generation via Physical Causal Interventions

The fundamental cause of OOD generalization failure in deep single-shot fringe demodulation is that single-pixel underdetermination encourages the network to fit spurious statistical correlations among workpiece surface reflectance, spatial illumination distributions, and phase. According to the forward structured-light image-formation model, when the camera–projector relative pose and object geometry are fixed, the geometric phase field ϕ is determined by the system's spatial transfer mapping; material reflectance properties,

ambient illumination, and sensor exposure gain act only on the photometric parameter set $\eta = \{A, B\}$.

We therefore formulated a physical causal intervention mechanism. While strictly fixing the spatial pose of the measured object and the projected pattern, we actively manipulate imaging photometric parameters to generate cross-domain pairs of fringe observations with identical geometric phase but different photometric appearances. This data organization combines two complementary mechanisms: offline paired acquisition across photometric domains and online synthesis of continuous physical perturbations.

1. Offline pairing across multiple photometric domains: With the object's pose and the calibrated structured-light system held constant, imaging parameters such as camera exposure time, ambient illumination intensity, and projected intensity gain are systematically varied. Fringe sequences acquired under different photometric conditions are then organized into offline observation pairs with strict pixel-level spatial alignment.
2. Online synthesis of continuous physical perturbations: To bridge the gaps between discretely sampled offline domains, the training pipeline attempts online forward photometric synthesis with a probability of 50% for the photometrically perturbed member of each triplet, using the source-domain background intensity A, modulation amplitude B, and unit complex phasor p. Here, $p = e^{i\phi} = \cos\phi + i\sin\phi$, as defined in Equation (4), and thus $\mathrm{Re}(p) = \cos\phi$. Therefore, (A, B, p) provides an equivalent representation of (A, B, ϕ) for the sinusoidal synthesis in Equation (7). If online synthesis is not selected, or if its validity checks fail, the original offline cross-domain paired sample is retained.

$$\tilde{I}_\eta = Q_8\left(\mathrm{clip}_{[0,1]}\left[A_\eta + B_\eta \cos\phi + n_\eta\right]\right) \tag{7}$$

where $Q_8(z) = \mathrm{round}(255z)/255$ denotes 8-bit quantization of normalized intensity, and $\mathrm{clip}_{[0,1]}$ restricts the intensity to the valid range. The term n_η represents the combined Poisson shot noise and Gaussian readout noise. No explicit nonlinear camera-response function or inverse-response correction is employed. The photometric components are modeled in the normalized intensity domain, assuming an approximately linear response within the nonsaturated valid region. According to the forward irradiance model, the perturbation parameters are generated using the following linear relations:

$$A_\eta = L_{\mathrm{amb},\eta} + \rho_\eta L_{\mathrm{dc},\eta}, \qquad B_\eta = \rho_\eta L_{\mathrm{ac},\eta} v_\eta \tag{8}$$

Here, ρ_η is a randomized surface-reflectance gain factor, and $v_\eta \in [0, 1]$ is a visibility factor describing occlusion by projected shadows.

The online photometric augmentation samples a multiplicative reflectance gain from $U(0.7, 1.3)$, and a spatial gain field by independently sampling a 6×6 grid from $U(0.85, 1.15)$ and bilinearly interpolating it to image resolution. The exposure gain is 2^u, with $u \sim U(-0.5, 0.5)$ (i.e., -0.5 to $+0.5$ EV), and the additional ambient component is sampled from $U(0, 0.015)$ in normalized intensity units.

Applied to the original fields A and B, these factors generate new photometric fields while keeping the phase label and pixel coordinates fixed; individual projector DC/AC terms and visibility are not independently randomized in this implementation. Poisson shot noise is simulated with an electron-count scale of 10,000 and independent zero-mean Gaussian readout noise of standard deviation 2/10,000, followed by 8-bit requantization.

An online candidate is accepted only if its modulation is at least 5 DN, its intensity is within the nonsaturated physical range, and its remaining valid-pixel coverage is at least 50% of the original mask. Otherwise, the offline photometric pair is retained.

To avoid mistaking physically irreversible degradation for algorithmic error, all counterfactual interventions are defined and evaluated only within the phase-observable mask Ω. The local valid-region mask is defined as

$$\Omega = \left\{ (x,y) \mid B_{\mathrm{gt}}(x,y) > B_{\mathrm{th}} \ \wedge \ I_k(x,y) < I_{\mathrm{sat}} \right\} \tag{9}$$

where $B_{\mathrm{th}} = 5$ DN is the minimum modulation threshold used to exclude regions where phase information is unreliable due to a low signal-to-noise ratio, and $I_{\mathrm{sat}} = 255$ DN is the saturation threshold used to exclude saturated sensor pixels. Since 8-bit grayscale images are used, the corresponding normalized thresholds are 5/255 and 1.0, respectively. The same phase-observable mask is consistently applied to all compared methods to ensure fair and unbiased evaluation.

2.4. Joint Learning of Phase Invariance and Equivariance

Using the data structure produced by the physical interventions, we impose two complementary physical symmetry constraints on the model outputs to encourage a feature space with physically correct causal responses.

2.4.1. Phase-Invariance Regularization Under Photometric Perturbation Conditions

Let T_η denote a photometric intervention that preserves object geometry and projected phase shift. Suppose that an input fringe is transformed with different photometric parameters into an anchor sample $I_a = T_{\eta_1} I$ and a photometrically paired sample $I_p = T_{\eta_2} I$. An ideal demodulation model $F_\theta(\cdot)$ should be strictly invariant to these photometric perturbations: $F_\theta(I_a) = F_\theta(I_p)$. Using the unit complex phasor as the metric representation, we constructed the periodic invariance loss over the overlapping valid region $\Omega_{ap} = \Omega_a \cap \Omega_p$:

$$\mathcal{L}_{\mathrm{inv}} = \frac{1}{|\Omega_{ap}|} \sum_{x \in \Omega_{ap}} \left[1 - \mathrm{Re}\left(\hat{p}_a(x) \overline{\hat{p}_p(x)} \right) \right] \tag{10}$$

where $\overline{\hat{p}}$ denotes complex conjugation. By the complex-phasor inner-product identity, $\mathrm{Re}(\hat{p}_a \overline{\hat{p}_p}) = \cos(\hat{\phi}_a - \hat{\phi}_p)$, so the loss attains its theoretical minimum of zero if and only if the predictions from the two domains have no relative phase difference. This constraint forces the network to suppress spurious dependence on material reflectance and illumination gradients.

2.4.2. Phase-Shift Equivariance Regularization with Projected Pattern Displacements

An invariance constraint alone risks producing constant predictions or excessive smoothing. To preserve sensitivity to genuine geometric phase variations, we introduced a projected phase-shift equivariance constraint. Applying a known phase shift δ to the projected grating (typically the quadrature shift $\delta = \pi/2$ during training) theoretically rotates the complex phasor by $R_\delta(\hat{p}_a) = \hat{p}_a \exp(i\delta)$. Given a pair of fringe images I_a and $I_s = I(d_1, \delta)$ of the same geometry but with different phase-shift steps, the network outputs must satisfy the rotational equivariance relation

$$\mathcal{L}_{\mathrm{eq}} = \frac{1}{|\Omega_{as}|} \sum_{x \in \Omega_{as}} \left[1 - \mathrm{Re}\left(\hat{p}_s(x) \overline{\hat{p}_a(x) e^{i\delta}} \right) \right] \tag{11}$$

Equation (11) requires the model to remain sensitive to phase-carrier deformations associated with surface height variations and ensures an appropriate response of the feature space to the phase-shifting operator.

2.4.3. Multi-Task, Multiscale Joint Optimization Objective

Combining task supervision, multiscale auxiliary supervision, and forward physical reconstruction, the overall training loss is defined as

$$\begin{aligned}\mathcal{L}_{\text{total}} = \lambda_{\text{phase}}\mathcal{L}_{\text{phase}} + \lambda_{\text{photo}}\mathcal{L}_{\text{photo}} + \lambda_{\text{render}}\mathcal{L}_{\text{render}} \\ + \lambda_{\text{grad}}\mathcal{L}_{\text{grad}} + \lambda_{\text{aux}}\mathcal{L}_{\text{aux}} + r(t)\left(\lambda_{\text{inv}}\mathcal{L}_{\text{inv}} + \lambda_{\text{eq}}\mathcal{L}_{\text{eq}}\right)\end{aligned} \tag{12}$$

Here, $\lambda_{phase} = 1.0$, $\lambda_{photo} = \lambda_{render} = \lambda_{grad} = \lambda_{aux} = 0.1$, and $\lambda_{inv} = \lambda_{eq} = 0.2$. The term $r(t) \in [0, 1]$ denotes a curriculum-based warm-up factor designed for the physical symmetry-consistency losses. It is set to zero during epochs 1∼5 to stabilize the initial phasor representation, then increases linearly, reaching its full weight of 1.0 at epoch 15. The component losses are specified below:

- Primary phase supervision and multiscale auxiliary loss: A periodic metric is defined using the high-accuracy reference phasor p_{gt} obtained by multi-step phase shifting:

$$\mathcal{L}_{\text{phase}} = \frac{1}{|\Omega|}\sum_{x\in\Omega}\left[1 - \text{Re}\left(\hat{p}(x)\overline{p_{\text{gt}}(x)}\right)\right] \tag{13}$$

  The multiscale auxiliary loss $\mathcal{L}_{\text{aux}}$ applies the same cosine-inner-product supervision to downsampled reference phasors at decoder levels $P2^{\dagger}$ and $P1^{\dagger}$.
- Photometric decoupling supervision: The auxiliary branch is encouraged to estimate the input domain's background intensity A_{gt} and modulation field B_{gt} (or the synthesis parameters A_η, B_η for online-perturbed samples):

$$\mathcal{L}_{\text{photo}} = \frac{1}{|\Omega|}\sum_{x\in\Omega}\left(|\hat{A}(x) - A_{\text{gt}}(x)| + |\hat{B}(x) - B_{\text{gt}}(x)|\right) \tag{14}$$

- Forward physical reconstruction loss: The decoded variables reconstruct the observed fringe intensity as $\hat{I} = \hat{A} + \hat{B}\hat{c}$. A robust Charbonnier penalty suppresses outlying noise:

$$\mathcal{L}_{\text{render}} = \frac{1}{|\Omega|}\sum_{x\in\Omega}\rho_{\text{Ch}}\left(\hat{A}(x) + \hat{B}(x)\,\text{Re}(\hat{p}(x)) - I(x)\right) \tag{15}$$

  where $\rho_{\text{Ch}}(z) = \sqrt{z^2 + \gamma^2}$ and $\gamma = 10^{-3}$ is a smoothing constant.
- Spatial phase-gradient continuity loss: First-order differential consistency of the local surface geometry is enforced by

$$\mathcal{L}_{\text{grad}} = \frac{1}{|\Omega_g|}\sum_{(x,\Delta)\in\Omega_g}\left|\hat{p}(x+\Delta)\overline{\hat{p}(x)} - p_{\text{gt}}(x+\Delta)\overline{p_{\text{gt}}(x)}\right|^2 \tag{16}$$

  where $\Delta \in \{(0, 1), (1, 0)\}$ denotes the horizontal and vertical nearest-neighbor offsets, respectively, and $\Omega_g = \{(x, \Delta) \mid x \in \Omega \wedge x + \Delta \in \Omega\}$ is the set of pixel pairs for which both neighbors are valid.

2.5. Triplet Sampling and Collaborative Training Strategy

The overall workflow of PCIE-FPP is summarized in Figure 4a. During training, multi-frequency fringe observations and their reference labels are used to construct physically intervened triplets, which are processed by the shared dual-branch network under supervised phase–photometric learning and joint invariance–equivariance constraints. During inference, the trained network recovers the wrapped phase from a single fringe image and its frequency index. The recovered phase maps at the three frequencies are subsequently used for multi-frequency unwrapping and calibrated 3D reconstruction. The detailed triplet construction, shared-network symmetry constraints, joint training objective, and single-shot inference procedures are illustrated in Figure 4b–e.

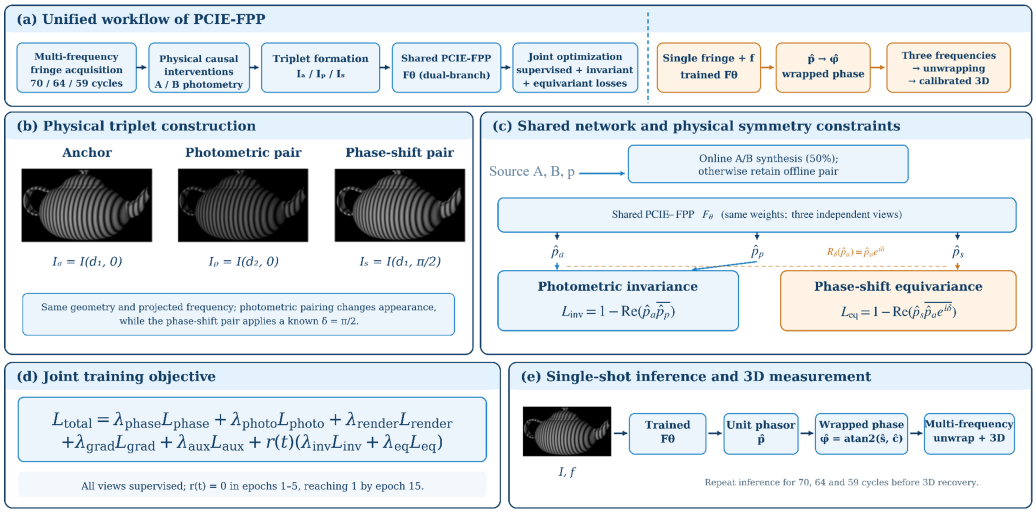


Figure 4. Unified workflow of the proposed PCIE-FPP framework for physical causal intervention, joint phase invariance–equivariance learning, and single-shot phase recovery. (**a**) Overall training and inference workflow, from multi-frequency fringe acquisition and reference-label generation to physical intervention, dual-branch network optimization, wrapped-phase recovery, and subsequent 3D reconstruction. (**b**) Construction of physical triplets comprising a source-domain anchor $I_a = I(d_1, 0)$, a photometrically perturbed sample $I_p = I(d_2, 0)$, and a phase-shifted sample $I_s = I(d_1, \pi/2)$. Online photometric synthesis is attempted with a probability of 50%; otherwise, or if the candidate fails validity checks, the offline photometric pair is retained. (**c**) Weight-shared collaborative training of the three views and the associated phase-invariance and phase-shift-equivariance constraints. (**d**) Joint training objective combining supervised phase–photometric losses, physical reconstruction, and the invariance–equivariance regularization terms. (**e**) Single-shot inference and 3D measurement, in which a fringe image and its frequency index are processed by the trained network to recover the wrapped phase; the three frequency-specific wrapped-phase maps are then used for multi-frequency unwrapping and calibrated 3D reconstruction.

To activate phase-invariance and phase-shift equivariance constraints simultaneously, we employ a physical-triplet data-loading pipeline during training (Figure 4). For each training batch, a triplet $\mathcal{T} = \{I_a, I_p, I_s\}$ is formed from the same object instance and projected frequency:

- $I_a = I(d_1, 0)$: The anchor fringe, from source photometric domain d_1 at phase-shift step 0.
- $I_p = I(d_2, 0)$: A counterfactual sample with the same geometry and phase shift as the anchor. An offline cross-domain pair is sampled first, and its photometric member is then selected for online synthesis with 50% probability; if no online candidate is selected or its validity checks fail, the offline pair is retained.
- $I_s = I(d_1, \pi/2)$: A rotationally equivariant sample from the same photometric domain d_1 as the anchor but with phase shift $\delta = \pi/2$.

Each mini-batch contains 10 physical triplets. Thus, each triplet comprises three single-frame views, and a mini-batch contains 30 single-frame views in total. The three views are passed separately through the weight-sharing PCIE-FPP network F_θ. Following the forward pass, $\mathcal{L}_{\text{inv}}$ is computed from the pair (I_a, I_p) and $\mathcal{L}_{\text{eq}}$ from (I_a, I_s). Each view independently receives full supervision from $\mathcal{L}_{\text{phase}}$, $\mathcal{L}_{\text{photo}}$, $\mathcal{L}_{\text{render}}$, $\mathcal{L}_{\text{grad}}$, and $\mathcal{L}_{\text{aux}}$ within its own valid mask. For the shifted view I_s, the reference-phasor label is rotated correspondingly to $p_{\text{gt}} e^{i\delta}$.

2.6. Experimental Setup and Evaluation Protocol

2.6.1. Optical Acquisition System and Multi-Domain Data Organization

As shown in Figure 5a, the experimental acquisition system consisted of a digital light processing (DLP) projection module and a monochrome industrial area-scan CMOS camera. The intrinsic matrices of the camera and projector and their stereo relative pose were calibrated in advance. During data acquisition, the projector sequentially displayed three-frequency, four-step phase-shifted sinusoidal patterns with fringe cycle counts $f_1 = 70$, $f_2 = 64$, and $f_3 = 59$ (corresponding to frequency indices $f \in \{0, 1, 2\}$). The measured objects included sculptural free-form surfaces, precision helical gear shafts, finned rack housings, and standard teapots, representing a range of workpieces and geometric objects with complex continuous curvatures, geometric steps, and sharp edges. Multiple views were acquired at different object orientations to form a comprehensive set of geometric surface observations.

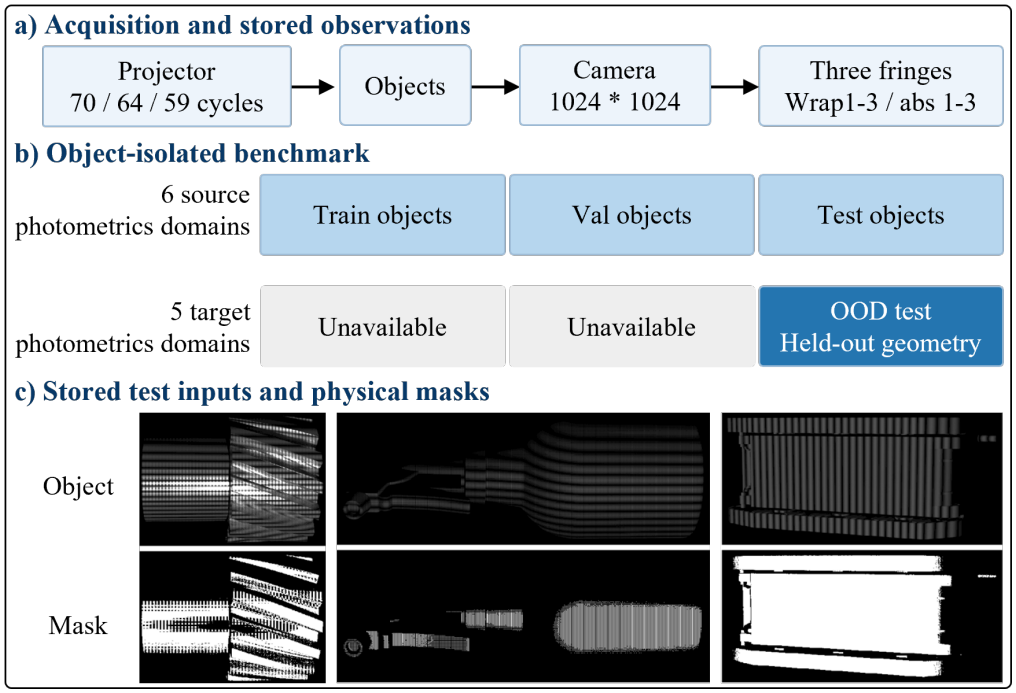


Figure 5. Fringe structured-light data acquisition and organization, and the dual-axis-isolated domain-generalization benchmark. (**a**) Data acquisition using three sinusoidal fringe frequencies with 70, 64, and 59 cycles, respectively, and four phase-shifting steps per frequency. The camera resolution is 1024 × 1024 pixels. Wrap 1–3 and abs 1–3 denote the corresponding wrapped-phase and absolute-phase maps. (**b**) Dataset partitioning independently along the object-instance and photometric-domain axes. The training, validation, and test sets contained 228, 29, and 28 nonoverlapping object instances, respectively. S1–S6 denote the six seen source photometric domains, whereas T1–T5 denote the five held-out target domains used for out-of-distribution (OOD) evaluation on unseen test objects. Target-domain observations were excluded from training and validation. (**c**) Representative complex industrial test workpieces and their phase-observable masks. The masks exclude pixels with insufficient fringe modulation ($B_{gt} \leq 5$ DN) or sensor saturation ($I_k \geq 255$ DN), where DN denotes digital number.

While strictly maintaining the pose of each measured object and all system calibration parameters, we systematically varied camera exposure time, ambient background illumination, projected fringe-intensity gain, and surface reflectance characteristics. The acquired fringe sequences were organized into six seen source photometric domains and five unseen target photometric domains (Figure 5b). Complex workpiece scenes exhibiting severe

diffuse-reflectance attenuation, localized specular highlights, and interreflections within deep grooves (Figure 5c) may deviate locally from the ideal sinusoidal model. These were included among the unseen target challenge domains to test generalization near nonideal physical boundaries.

The relative photometric settings of the six source domains (S1–S6) and five held-out target domains (T1–T5) are summarized in Table 1. The baseline S1 defines zero exposure offset and unit projector and ambient-illumination factors. Exposure offsets are expressed in EV; projector gain and ambient-illumination factors are dimensionless quantities relative to S1. The five target photometric configurations were excluded from training and hyperparameter selection.

Table 1. Relative photometric configurations of the six source domains and five held-out target domains.

Domain ID	Exposure Offset (EV)	Projector Gain	Ambient Factor	Domain-Specific Condition
S1 Baseline	0.0	1.0	1.0	Standard diffuse surface
S2 Low Exposure	−0.5	1.0	1.0	Low-exposure diffuse surface
S3 High Exposure	+0.5	1.0	1.0	High-exposure diffuse surface
S4 Dark Surface	0.0	1.0	1.0	Dark diffuse surface
S5 Light Surface	0.0	1.0	1.0	Light-colored diffuse surface
S6 Spatial Variation	0.0	1.0	1.0	Smooth spatial reflectance variation
T1 Extremely Low Exposure	−2.0	1.0	1.0	Extremely low-exposure diffuse surface
T2 Weak Projected Illumination	0.0	0.25	2.0	Reduced projected fringe illumination and increased ambient light
T3 Bright Metallic Surface	0.0	1.0	1.0	Bright metallic surface
T4 Dark Metallic Surface	0.0	1.0	1.0	Dark metallic surface
T5 Spatial Variation	0.0	1.0	1.0	Target-domain spatial reflectance variation

2.6.2. Reference Ground Truth and Dataset Partitioning

To obtain accurate supervision labels and evaluation references, complete multi-step, three-frequency phase-shifting fringe sequences were acquired for each object at a fixed pose under nominal exposure and standard illumination conditions. Reference data were then constructed as follows. Standard multi-step quadrature demodulation extracted pixelwise reference wrapped phases at the three frequencies (Wrap 1–3) and photometric reference fields (A_{gt}, B_{gt}) from the nominal sequences. Multi-frequency heterodyne unwrapping then recovered the full-field reference absolute phases (abs 1–3), and the calibrated system parameters were used to compute the reference 3D depth map Z_{gt}. Because the pose of an object remains strictly unchanged when photometric acquisition conditions are switched, single-frame fringe inputs from different photometric domains are naturally aligned at the pixel level with the reference phase labels and depth map recovered from the corresponding nominal sequence.

As illustrated in Figure 5b, the dataset was rigorously divided along two independent axes, geometry instance × imaging photometric domain. On the geometry axis, each independent 3D object instance forms an isolated unit, ensuring that object topologies do not overlap among the training, validation, and geometry-extrapolation test sets. On the photometric axis, six domains were designated as seen source domains and five as unseen target domains. The unseen target photometric domains were completely excluded from network training and hyperparameter selection to prevent target-domain information leakage. At the object-instance level, the dataset was divided into training, validation, and test subsets at an approximate ratio of 8:1:1, comprising 228, 29, and 28 nonoverlapping objects, respectively. These subsets correspond to 912, 116, and 112 geometry-view records, respectively.

Only source-domain observations from the training and validation subsets were used for model fitting and validation, whereas the 28 held-out test objects were evaluated under both source-domain and target-domain photometric conditions.

2.6.3. Training Configuration and Optimization

The proposed PCIE-FPP model was trained using the AdamW optimizer with an initial learning rate of 2×10^{-4}, weight decay of 1×10^{-4}, and a batch size of 10 physical triplets (30 single-frame views per mini-batch). The training schedule consisted of 300 epochs, with a five-epoch learning-rate warm-up followed by cosine annealing. The phase-invariance and phase-shift-equivariance regularization terms were disabled for the first five epochs and linearly ramped to their full weights by epoch 15, as specified in Section 2.4. The seven loss weights are explicitly given immediately after Equation (12); the model parameter count is reported in the "Computational Complexity and Inference Timing" subsection. The common cross-method training and evaluation protocol is described in Section 2.6.4.

2.6.4. Comparison Baselines and Quantitative Metrics

Two complementary evaluation settings were adopted for a comprehensive performance assessment. (1) The cross-method comparison includes three representative deep-learning fringe-demodulation and domain-adaptation approaches. Method 1 denotes a single-shot approach trained using deep learning and computer-graphics synthesis [30]; Method 2 denotes a hybrid-encoding, simulation-to-real transfer approach [31]; and Method 3 denotes a structured-light reconstruction approach combining spatial–frequency information with domain-adaptive adversarial learning [32]. (2) The ablation study examines conventional quadrature-component (N/D) regression, the unit complex phasor representation, online physical augmentation, invariance/equivariance constraints, and photometric learning components in configurations A0–A6 and B0–B2.

Baseline training and evaluation protocol. To ensure a statistically comparable cross-method comparison, all four methods (Method 1, Method 2, Method 3, and PCIE-FPP) were trained and validated using the same object-level data partitions and the same six source photometric domains (S1–S6). The training, validation, and test subsets contained 228, 29, and 28 nonoverlapping object instances, respectively. Only source-domain observations were used for training and validation; the five target photometric domains (T1–T5) were excluded from training and hyperparameter selection. The 28 held-out test objects were not used for model fitting or selection.

All four methods were evaluated on the same 28 held-out test objects under both seen source-domain and unseen target-domain photometric conditions. The same reference wrapped-phase maps and phase-observable masks were used for every method. Circular phase MAE and RMSE were calculated using the same error definition, and depth accuracy was assessed using the same multi-frequency phase-unwrapping and calibrated 3D reconstruction procedure. The reported results were aggregated as object-level macro-averages over the same test objects. To accommodate differences in the original network architectures and fringe-encoding requirements, method-specific input representations and optimization hyperparameters were retained where necessary. Thus, comparability was controlled through the data partitions, domain exposure, test objects, reference labels, masks, and evaluation pipeline, rather than by forcing different architectures to use identical optimizer settings.

The evaluation uses a circular phase-error metric:

$$e_{\phi}(x) = \operatorname{atan2}\Big(\sin[\hat{\phi}(x) - \phi_{\mathrm{gt}}(x)],\ \cos[\hat{\phi}(x) - \phi_{\mathrm{gt}}(x)]\Big) \tag{17}$$

For unseen test objects, object-level macro-averages were calculated for wrapped-phase mean absolute error (MAE) and root-mean-square error (RMSE) in seen and unseen photometric domains, as well as the mean absolute 3D depth error after phase unwrapping and the cross-sectional residual ΔZ.

3. Results

This section reports the comparative wrapped-phase errors, continuous and compound photometric-domain sweeps, qualitative phase and depth residuals, physical-response diagnostics, ablation measurements, and three-dimensional reconstruction results on held-out test objects. Interpretations in relation to previous studies, underlying mechanisms, practical significance, and limitations are provided separately in Section 4.

3.1. Generalization Results Across Photometric Domains

3.1.1. Comparison Across Seen and Unseen Photometric Domains

Table 2 summarizes the wrapped-phase recovery errors of the four methods on unseen test objects, evaluated separately in seen and unseen photometric domains.

Table 2. Wrapped-phase errors in seen and unseen photometric domains (rad).

Method	Seen-Domain MAE (rad)	Seen-Domain RMSE (rad)	Unseen-Domain MAE (rad)	Unseen-Domain RMSE (rad)
Method 1	0.0868	0.1775	0.1764	0.3456
Method 2	0.0734	0.1532	0.1546	0.3253
Method 3	0.0617	0.1307	0.1426	0.3140
PCIE-FPP (Ours)	0.0552	0.1157	0.0824	0.2066

Table 2 shows that all four methods have larger phase errors in unseen than in seen photometric domains. Method 1 has seen/unseen MAEs of 0.0868/0.1764 rad, Method 2 0.0734/0.1546 rad, and Method 3 0.0617/0.1426 rad. PCIE-FPP has the lowest seen/unseen MAEs of 0.0552/0.0824 rad and the lowest seen/unseen RMSEs of 0.1157/0.2066 rad. Relative to Method 3, its unseen-domain MAE and RMSE are lower by approximately 42.2% and 34.2%, respectively.

3.1.2. Continuous Single-Parameter Domain Sweeps

To characterize the boundaries of performance degradation outside the training distribution, we performed continuous extrapolation sweeps along individual physical parameters while fixing object geometry, pose, and all other system parameters. Specifically, we examined changes in the fringe-modulation scaling factor, additional ambient stray-light intensity, and camera exposure offset.

Figure 6a–c present the circular-phase MAE curves (with logarithmic vertical axes) and valid-pixel coverage ratios for the four models over the continuous scan ranges. The results show the following. (1) During the modulation sweep (Figure 6a), as β decreases from its nominal value of 1.0 to the weak-modulation interval $[0.2, 0.5]$, phase errors of Method 1 and Method 2 rapidly exceed the 0.5∼1.0 rad range. PCIE-FPP (solid curve) maintains the lowest error envelope until $\beta > 1.5$ causes extensive sensor saturation clipping and a sharp drop in the valid-pixel ratio, at which point all methods approach physical limits. (2) During the ambient stray-light sweep (Figure 6b), as additional illumination increases from 0 DN to 80 DN, the error of PCIE-FPP increases at a substantially lower rate than those of the baselines. Even under the extreme condition DN = 80, its error remains below 0.40 rad, compared with 0.67∼0.89 rad for the baselines. (3) During the camera exposure sweep (Figure 6c), PCIE-FPP exhibits a broad, nearly flat, basin-like response over $[-2.5, +1.5]$ EV. Over this interval, its measured phase-error curve varies less than those of the compared methods.

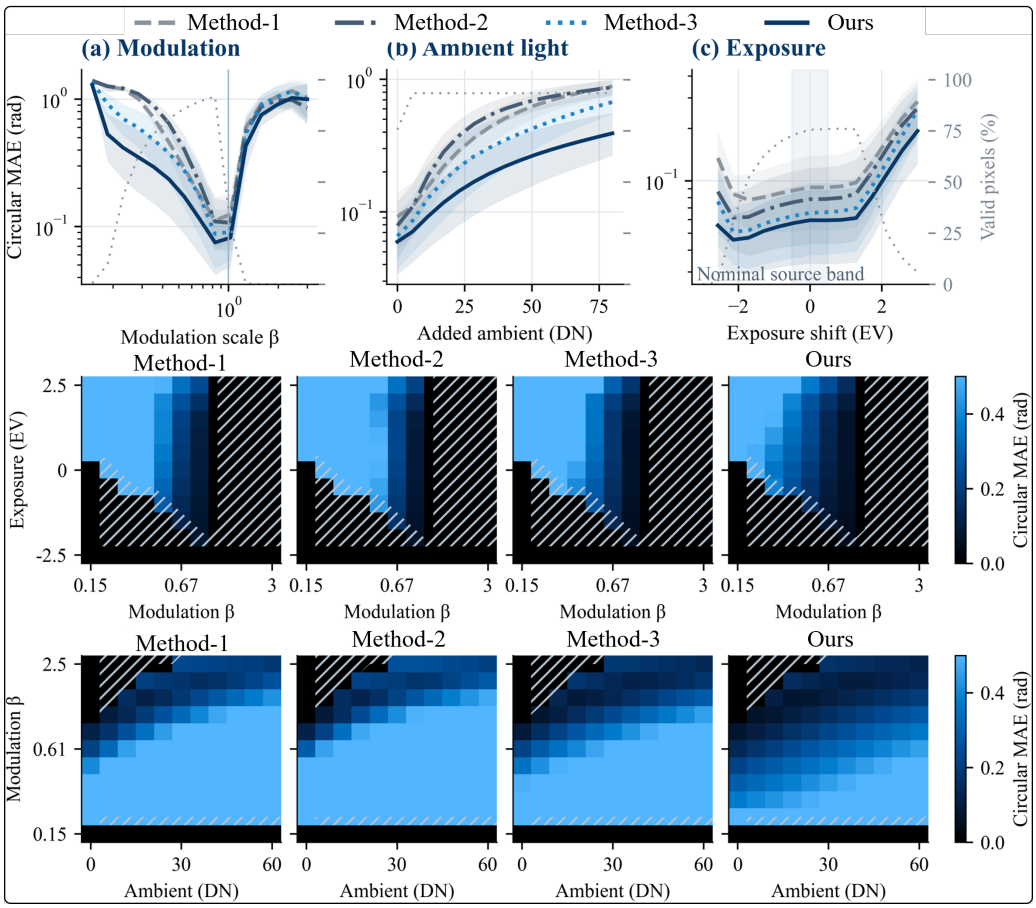


Figure 6. Continuous single-parameter extrapolation curves and two-parameter compound-domain-shift error maps. (**a**–**c**) Single-parameter sweeps of the fringe-modulation scaling factor β (dimensionless), added ambient illumination (DN, digital number), and camera exposure offset (EV, exposure values), respectively, with the remaining imaging parameters fixed. The primary evaluation metric is circular wrapped-phase mean absolute error (MAE, rad), plotted on a logarithmic vertical scale. Valid-pixel percentages are indicated using the secondary vertical axis. The second and third rows show compound-domain circular-phase MAE heatmaps over the parameter spaces "exposure offset (EV) × modulation scaling factor β" and "modulation scaling factor β × added ambient illumination (DN)", respectively. The heatmap color scales represent circular-phase MAE in radians and are shared across methods within each parameter grid.

3.1.3. Compound Two-Parameter Domain Extrapolation

In practical industrial inspection, several physical variables often drift simultaneously. The second and third rows of Figure 6 therefore consider two grids of compound perturbations: "camera exposure EV× modulation amplitude β" and "modulation amplitude β× ambient background intensity DN".

The two-dimensional error heatmaps show large bright-blue, high-error regions ($MAE_{circ} > 0.35$ rad) for Method 1 and Method 2 where both variables deteriorate in the same direction. Although Method 3 offers some improvement, it still degrades rapidly in the coupled low-modulation/high-ambient-light regime. In contrast, PCIE-FPP (the rightmost heatmaps) substantially expands the coverage of dark, low-error regions ($MAE_{circ} < 0.15$ rad) in both parameter spaces, as quantified by the corresponding circular-phase MAE color scales.

3.2. Wrapped-Phase Recovery Accuracy and Spatial Surface Fidelity

Figure 7 shows single-shot fringe inputs, recovered wrapped phases, wrapped-phase residual maps within the valid mask Ω_ϕ, and the propagated absolute-depth error distributions for two representative, previously unseen complex industrial workpieces under unknown photometric conditions.

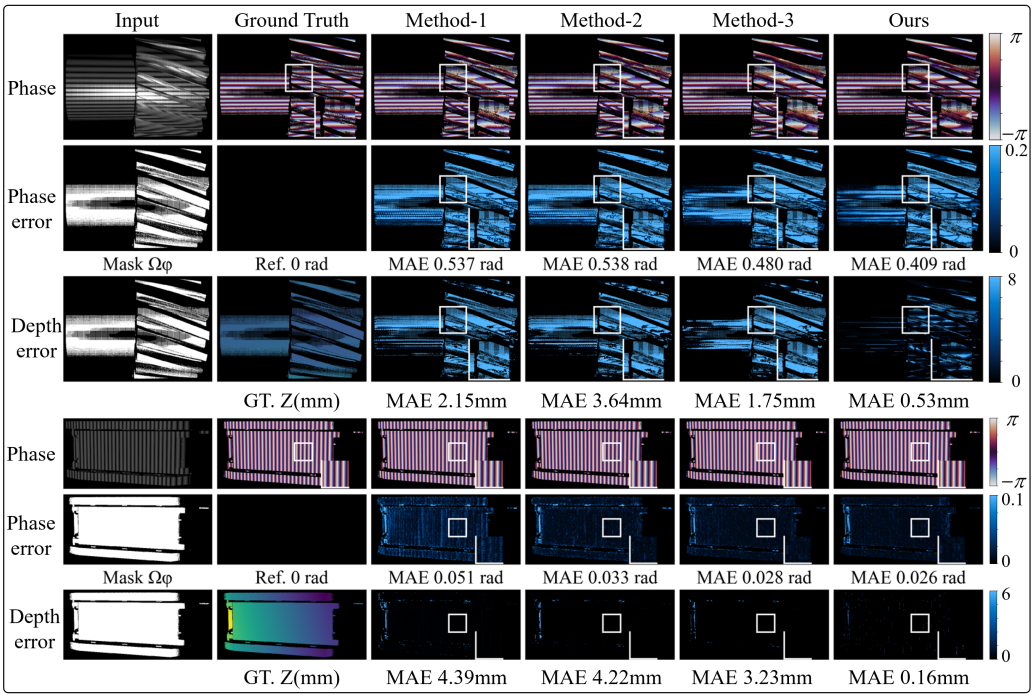


Figure 7. Qualitative and quantitative comparisons of wrapped-phase recovery and propagated depth errors on unseen complex industrial workpieces. The upper three rows show a highly reflective complex helical gear shaft, and the lower three rows show a rack. From left to right: single-shot input fringe and physical mask Ω_ϕ, reference ground truth, Method 1, Method 2, Method 3, and the proposed PCIE-FPP. White rectangles highlight enlarged regions near gear-tooth edges and fine grooves. The phase maps, absolute phase-error maps, and absolute depth-error maps use color bars in radians (rad), radians (rad), and millimeters (mm), respectively. Their numerical display ranges are shown on the color bars and are kept consistent across methods within each test scene.

The sample-specific measurements and residual distributions in Figure 7 reveal two key findings. (1) The left shaft surface and right helical gear surface exhibit abrupt reflectance changes and localized specular highlights. Method 1 and Method 2 yield wrapped-phase MAEs of 0.537 rad and 0.538 rad, respectively, with dense fringe-leakage ripples near helical tooth roots and shaft transitions. As a result, their reconstruction depth MAEs reach 2.15 mm and 3.64 mm. Method 3 reduces the phase and depth errors to 0.480 rad and 1.75 mm, whereas PCIE-FPP further lowers the wrapped-phase error to 0.409 rad, suppresses local tooth-surface phase distortions, and achieves a depth MAE of 0.53 mm. (2) In the presence of densely spaced, periodic vertical deep grooves and shadow boundaries, Method 1, Method 2, and Method 3 exhibit relatively small full-field average wrapped-phase errors of 0.051 rad, 0.033 rad, and 0.028 rad, respectively. However, local phase-error spikes at groove edges trigger fringe-order jumps during multi-frequency heterodyne unwrapping, increasing their depth MAEs to 4.39 mm, 4.22 mm, and 3.23 mm. By comparison, PCIE-FPP achieves the lowest wrapped-phase error (0.026 rad) and preserves edge gradients sufficiently well to limit the absolute-depth MAE to 0.16 mm.

3.3. Physical-Response Diagnostic Results

3.3.1. Counterfactual Consistency Under Identical Geometry and Different Photometric Conditions

Multiple fringe images were collected under varying photometric conditions with geometry and pose held fixed, and the composite magnitude variance $V(x)$ of predicted unit complex phasors across photometric domains was calculated. The theoretical minimum is $V(x) = 0$ when predicted phases agree exactly for the same pixel across all tested conditions. Figure 8a reports full-field mean variances of 2.16 $\times$ 10^{-4} for Method 1, 1.67×10^{-4} for Method 2, 6.56×10^{-5} for Method 3, and 7.75×10^{-5} for PCIE-FPP. PCIE-FPP reduces the variance relative to Method 1 and Method 2 by approximately 64.1% and 53.6%, respectively, while Method 3 has the lowest mean variance among the four models.

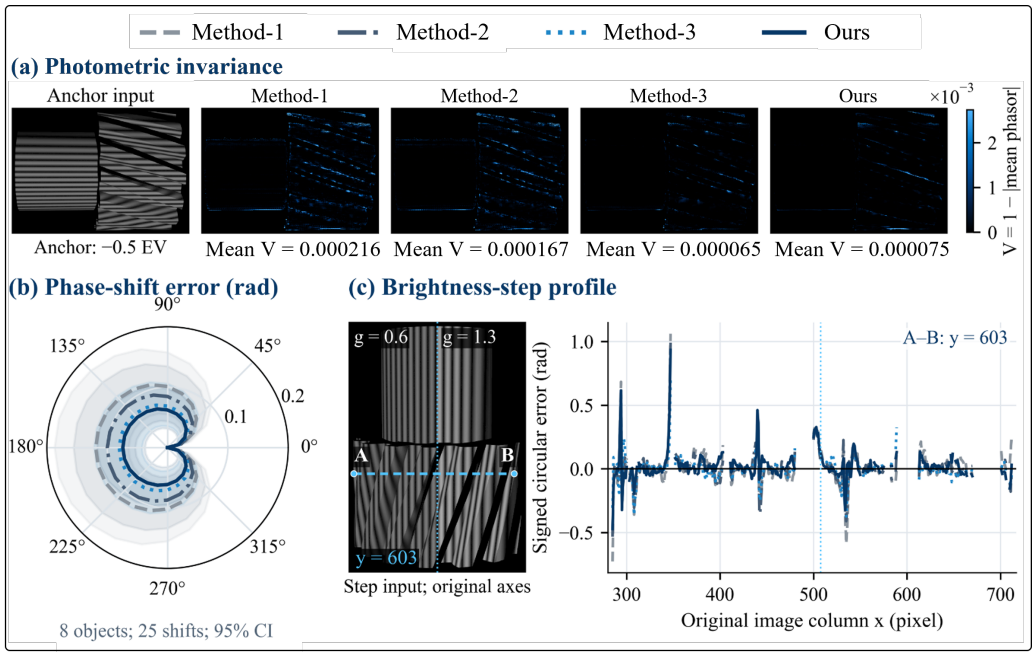


Figure 8. Diagnostics of photometric phase invariance, phase-shift rotational equivariance, and robustness to abrupt reflectance variations. (**a**) Counterfactual unit-phasor variance maps across photometric domains, using −0.5 EV as the anchor condition. The dimensionless metric V measures the dispersion of predicted unit complex phasors with photometric changes; lower values indicate stronger phase invariance. All compared methods use the same color scale, with the scaling factor indicated on the colorbar. (**b**) Polar distributions of phase-shift equivariance errors, evaluated on eight unseen test objects and 25 uniformly spaced projected phase shifts. The angular coordinate represents the applied phase shift, while the radial coordinate represents the circular equivariance error in radians. Shaded regions indicate 95% confidence intervals (CI). (**c**) A two-sided reflectance-step experiment, with gain $g = 0.6$ in the left half and $g = 1.3$ in the right half. Signed circular-phase errors (rad) are extracted along the horizontal cross-section A–B at $y = 603$ in the original 1024 × 1024 image coordinate system (zero-based indexing), corresponding to the same row in the full-resolution predicted phase map. The horizontal coordinate x is expressed in original-image pixels.

3.3.2. Rotational Equivariance with Identical Geometry and Different Phase Shifts

The phase-shift equivariance test is intended to show that the learned invariance does not degenerate into a trivial insensitivity to changes in the input waveform. All photometric and geometric scene parameters were kept fixed while 25 discrete projected phase shifts δ spanning $[0°, 360°]$ were applied, including arbitrary non-quadrature shifts not encountered during training. The circular equivariance deviation $E_{\mathrm{eq}}(\delta)$ was evaluated on eight independent test objects. The polar plots in Figure 8b show that Method 1 and

Method 2 have phase-shift equivariance errors as high as 0.12∼0.16 rad throughout the phase-shift cycle, with wide 95% confidence intervals. Method 3 reduces the error to 0.095 rad but retains off-center expansion in the non-quadrature phase-shift quadrants. In contrast, PCIE-FPP (dark-blue solid curve) remains close to the inner ring (0.082 rad) for all 25 shifts, exhibiting accurate, approximately isotropic rotational equivariance. The corresponding polar response is more uniform than those of the three compared methods across the tested shifts.

3.3.3. Nonuniform Reflectance-Step Boundaries and Complex-Phasor Orthogonality Diagnostics

To further examine the model's resistance to spatially abrupt reflectance changes, Figure 8c presents a challenging two-sided brightness-step experiment. A reflectance-gain step was imposed across a gear workpiece at the horizontal midpoint of the image, with gain $g = 0.6$ in the left half of the field of view and $g = 1.3$ in the right half. Signed circular-phase error profiles were extracted along the horizontal scan line A–B at y = 603 in the original 1024 × 1024 image coordinate system (zero-based indexing). Since the network output is restored to the original spatial resolution, this scan line corresponds directly to y = 603 in the predicted phase-error map, without additional coordinate transformation. The line crosses the brightness boundary (dashed line at x ≈ 505). The baseline methods show pronounced oscillations ranging from −0.6 to +0.3 rad near the step boundary and the high-gain gear teeth on the right. By contrast, PCIE-FPP maintains a stable, approximately zero-mean transition across the $g = 0.6 \rightarrow 1.3$ step and shows smaller error excursions near the low-modulation tooth-root regions.

3.4. Ablation Results

To isolate the contributions of phase-invariance, phase-equivariance, and photometric learning components, ten ablation configurations (A0–A6 and B0–B2) were evaluated with identical experimental settings. A2 served as the common reference for independently introducing $\mathcal{L}_{\text{inv}}$ and $\mathcal{L}_{\text{eq}}$ (A3 and A4), while A6 served as the common reference for independently removing $\mathcal{L}_{\text{photo}}$ and $\mathcal{L}_{\text{render}}$ (B1 and B2). Table 3 summarizes the corresponding phase and depth errors.

Table 3. Ablation results with different network components and physical constraints.

Configuration	Seen-Domain Phase MAE (rad)	Unseen-Domain Phase MAE (rad)	Seen-Domain Depth MAE (mm)	Unseen-Domain Depth MAE (mm)
A0	0.0780	0.1760	0.8200	1.7600
A1	0.0730	0.1650	0.7400	1.5700
A2	0.0660	0.1530	0.6500	1.3600
A3	0.0610	0.1410	0.5600	1.1200
A4	0.0630	0.1460	0.5900	1.1900
A5	0.0580	0.1370	0.4900	1.0200
A6	0.0552	0.0824	0.4453	0.9448
B0	0.0640	0.1450	0.6100	1.3150
B1	0.0568	0.1180	0.4730	1.0060
B2	0.0560	0.1060	0.4620	0.9840

Table 3 shows the progressive effects of the configurations. From quadrature-component regression to a unit complex phasor (A0 → A1): Configuration A0 directly regresses the conventional orthogonal components (N, D), which remain coupled to the

modulation amplitude. Its unseen-domain phase and depth MAEs are 0.1760 rad and 1.7600 mm, respectively. Configuration A1 introduces a manifold-normalized unit complex phasor to remove the scaling effect of the local photometric envelope B from the output representation and eliminate the periodic discontinuity at $\pm\pi$. The seen/unseen phase MAEs consequently fall to 0.0730/0.1650 rad, while the depth MAEs decrease to 0.7400/1.5700 mm.

Introducing online physics-based causal augmentation (A1 → A2): Configuration A2 adds online physical A/B photometric perturbations and paired multi-domain data augmentation to the unit-phasor baseline. This broadens the support of the source-domain photometric distribution and further reduces unseen-domain phase and depth MAEs to 0.1530 rad and 1.3600 mm.

Individual and joint phase-invariance/phase-shift-equivariance constraints (A2 → A3/A4 → A5): Configurations A3 and A4 independently add the cross-photometric-domain invariance constraint $\mathcal{L}_{\text{inv}}$ and the quadrature phase-shift equivariance constraint $\mathcal{L}_{\text{eq}}$ to A2. Their unseen-domain phase MAEs decrease to 0.1410 rad and 0.1460 rad, respectively. Configuration A5 applies both symmetry constraints simultaneously through physical triplets, achieving seen/unseen phase MAEs of 0.0580/0.1370 rad and depth MAEs of 0.4900/1.0200 mm. The associated unseen-domain depth MAE for A5 is 1.0200 mm.

Photometric-branch decoupling and the forward physical reconstruction loop (A5 → A6): The complete model A6 extends A5 with a photometric decoder branch (Q3–Q0), photometric decoupling supervision $\mathcal{L}_{\text{photo}}$, and a forward fringe reconstruction loss $\mathcal{L}_{\text{render}}$. A6 achieves the best performance on unseen test objects, with seen- and unseen-domain phase MAEs of 0.0552 rad and 0.0824 rad and 3D depth MAEs of 0.4453 mm and 0.9448 mm. Compared with the baseline A0, the respective depth errors are reduced by approximately 45.7% and 46.3%.

B0 adds the photometric branch and rendering loss to A2 without photometric supervision or symmetry constraints, while B1 and B2 remove the photometric supervision and rendering losses from A6, respectively. In the illustrative reference results, B0 achieves an unseen-domain phase MAE of 0.1450 rad, compared with 0.1530 rad for A2. Removing photometric supervision (B1) or rendering consistency (B2) increases the unseen-domain phase MAE from 0.0824 rad (A6) to 0.1180 rad and 0.1060 rad, respectively. The corresponding unseen-domain depth MAEs for B1 and B2 are 1.0060 and 0.9840 mm, respectively.

3.5. 3D Reconstruction Accuracy and Inference Timing

3.5.1. Multi-Frequency Phase Unwrapping and 3D Geometric Reconstruction Accuracy

To investigate how single-shot wrapped-phase estimates propagate to the final 3D reconstruction, we compile full-field 3D depth-error statistics for 28 unseen test objects in seen and unseen photometric domains (Table 4). For two representative unseen geometry instances, we also present color-coded 3D reconstructions under unknown photometric conditions, along with cross-sectional depth profiles and magnified residual curves extracted along reference line A–B (Figure 9).

Table 4. Depth MAE in seen and unseen photometric domains (mm).

Method	Seen-Domain MAE (mm)	Unseen-Domain MAE (mm)
Method 1	0.8130	3.8534
Method 2	0.5726	2.6989
Method 3	0.5413	2.3982
PCIE-FPP (Ours)	0.4453	0.9448

As shown in Table 4, when extrapolating from seen to unseen photometric domains, the depth MAEs of Method 1, Method 2, and Method 3 increase markedly from 0.8130,

0.5726, and 0.5413 mm to 3.8534, 2.6989, and 2.3982 mm, respectively. In contrast, PCIE-FPP yields depth MAEs of 0.4453 mm and 0.9448 mm in seen and unseen photometric domains, respectively, maintaining submillimeter reconstruction accuracy even under unknown photometric conditions. Compared with Method 1, Method 2, and Method 3, the proposed method reduces unseen-domain depth MAE by approximately 75.5%, 65.0%, and 60.6%, respectively (with corresponding seen-domain reductions of 45.2%, 22.2%, and 17.7%).

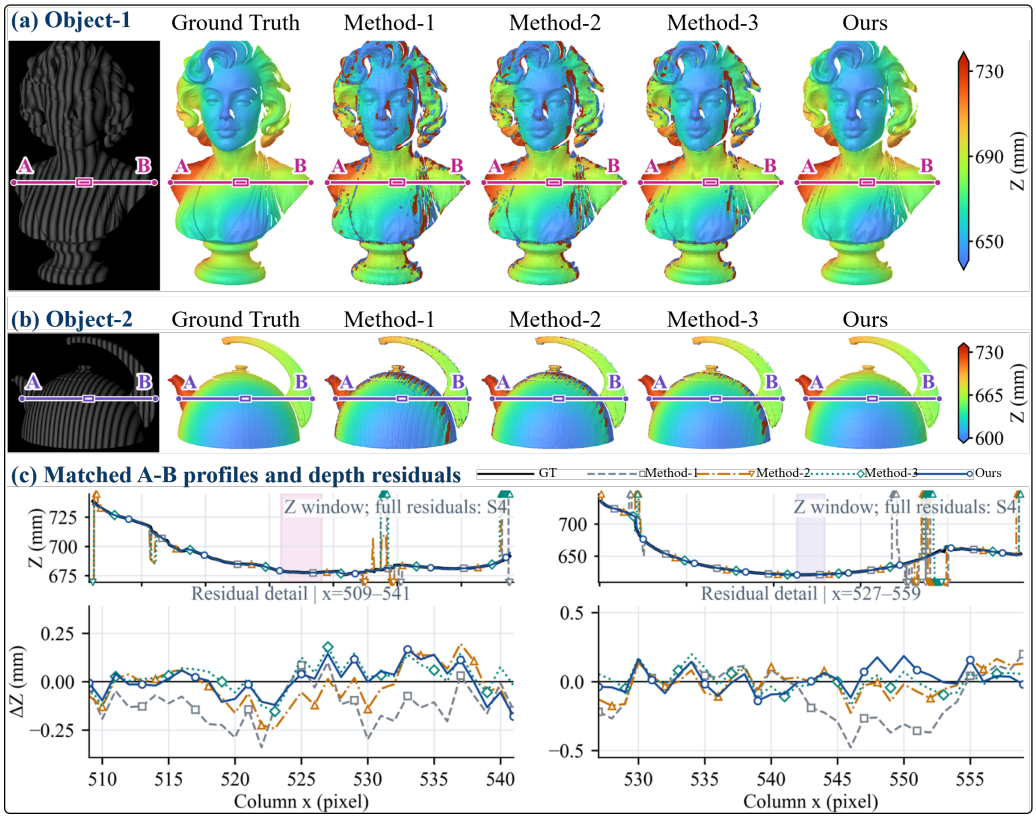


Figure 9. Multi-frequency heterodyne 3D surface reconstruction and cross-sectional depth-residual comparison on unseen geometry instances. (**a**) Object 1 and (**b**) Object 2: single-shot fringe inputs, ground truth, and color-coded 3D depth reconstructions using Method 1, Method 2, Method 3, and the proposed method (Ours) under unseen photometric conditions. (**c**) Absolute-depth profiles Z(mm) extracted along reference line A–B and depth-residual curves within a magnified local region. The color bars for the reconstructed depth maps are expressed in millimeters (mm), with their numerical ranges displayed alongside the reconstructions.

Figure 9 illustrates reconstruction details for two representative unseen geometry instances in unknown photometric domains. For Object-1, which has complex local curvature, the reconstructions of Method 1 and Method 2 show conspicuous step-like patches and local holes around recessed hair strands, nasal contours, and shadowed garment folds. Enlarged residual profiles along section A–B reveal a systematic negative bias of approximately −0.20∼−0.35 mm for Method 1, whereas PCIE-FPP maintains relatively stable sectional residuals within [−0.12, +0.16] mm. For Object-2, which has a smooth continuous surface, the baseline methods produce periodic depth ripples near the lid edge and around the highly reflective central region of the teapot body, with local deviations reaching −0.48 mm. In contrast, the reconstructed surface of PCIE-FPP closely matches the reference ground truth, demonstrating high geometric fidelity for both continuous free-form surfaces and complex fine-scale details.

3.5.2. Computational Complexity and Inference Timing

The proposed PCIE-FPP model contains 24.07 million trainable parameters. To evaluate its computational efficiency, inference experiments were conducted on a system equipped with an Intel Core i7-14700KF CPU (Intel Corporation, Santa Clara, CA, USA) and an NVIDIA GeForce RTX 5080 GPU (NVIDIA Corporation, Santa Clara, CA, USA), using PyTorch 2.7.1. For a single 1024 × 1024 grayscale fringe image, the measured inference time was 0.0282 s per frame (28.2 ms/frame).

4. Discussion

4.1. Relationship to Previous Fringe-Projection and Domain-Generalization Studies

The results indicate that the principal advantage of PCIE-FPP is not confined to accuracy in familiar imaging conditions; it is the smaller degradation when object appearance changes without a corresponding change in geometric phase. In Table 2, the phase MAE increases from 0.0552 to 0.0824 rad between seen and unseen photometric domains, an absolute increase of 0.0272 rad. The corresponding increase for Method 3 is 0.0809 rad (0.0617 to 0.1426 rad). Under the shared evaluation protocol, PCIE-FPP therefore reduces unseen-domain phase MAE by 42.2% relative to Method 3. These comparisons support improved resistance to the tested photometric shifts, but should not be interpreted as universal rankings across different acquisition systems or datasets.

This finding complements several established directions in learning-based fringe projection profilometry (FPP). Wang et al. used computer-graphics-based virtual FPP data to support single-shot prediction [30], whereas Ren et al. combined hybrid fringe encoding with simulation-to-real learning for depth estimation [31]. Jiang et al. investigated spatial–frequency feature fusion and domain-adversarial learning to improve cross-domain structured-light imaging [32]. These approaches address data availability, encoding, and distribution alignment. In comparison, PCIE-FPP explicitly specifies which image changes should leave phase unchanged and which projected-pattern changes should rotate the predicted phase. Consequently, its distinctive contribution is not simply wider appearance augmentation, but a physically defined response constraint that can be tested with counterfactual photometric pairs and controlled phase shifts. Because the compared approaches differ in their input encoding, targets, and model design, the reported numerical advantages are specific to the implementations and common test protocol described in Section 2.6.

Physics-informed fringe analysis has likewise demonstrated the value of embedding optical knowledge into learning. Yin et al. combined a learning-enhanced Fourier-transform module with a neural network [13], while Li et al. incorporated calibrated geometric constraints into single-shot three-dimensional measurement [14]. PCIE-FPP extends this line of research in a different direction: the unit-phasor representation separates geometric phase orientation from modulation magnitude, and the training objective constrains the predicted phase in two complementary, physically interpretable interventions. Importantly, the present three-dimensional evaluation still uses three-frequency phase unwrapping; it should therefore not be equated with a complete single-exposure absolute-depth reconstruction method.

4.2. Interpretation of the Invariance–Equivariance Mechanism and Ablation Evidence

The physical motivation for the method follows from the fringe model $I = A + B\cos\phi$. Under fixed geometry and approximately linear conditions, single-path imaging, changing A or positive B modifies fringe appearance without changing ϕ. The unit complex phasor retains phase direction while removing the modulation scale, and the differential relation in Equation (3) shows that phase error depends on the tangential component of quadrature error relative to B. This explains why amplitude normalization is relevant to photometric

robustness, while also showing why it cannot recover information once B becomes too small. In the ablation results (Table 3), replacing conventional quadrature-component regression (A0) with a unit phasor (A1) lowers unseen-domain phase MAE from 0.1760 to 0.1650 rad; adding online physical augmentation (A2) lowers it further to 0.1530 rad. Thus, the representation and broader training distribution contribute, but neither fully accounts for the complete model's performance.

The two symmetry constraints address different potential failure modes. Photometric invariance discourages a change in phase prediction when reflectance, illumination, or exposure changes at fixed geometry. Equivariance instead requires a known projected phase shift δ to produce the corresponding complex rotation $e^{i\delta}$, preventing invariance from collapsing into general insensitivity to the fringe signal. Starting from A2, adding invariance alone (A3) or equivariance alone (A4) gives unseen-domain phase MAEs of 0.1410 and 0.1460 rad, respectively; combining them (A5) gives 0.1370 rad. Incorporating the photometric branch and rendering-related objectives in A6 reduces the error further to 0.0824 rad. The B1 and B2 removals increase this error to 0.1180 and 0.1060 rad, respectively. These controlled comparisons are consistent with complementary contributions, although they do not establish that individual losses act independently or that every gain can be attributed to a single component.

An instructive qualification emerges from Figure 8. Method 3 has a slightly lower average cross-domain phasor variance (6.56×10^{-5}) than PCIE-FPP (7.75×10^{-5}), yet its reported phase-shift equivariance error is higher (0.095 versus approximately 0.082 rad). Therefore, minimizing photometric variation alone is not a sufficient criterion for reliable phase demodulation: a model must also respond correctly to real changes in phase. The data are consistent with the intended balance between these requirements. However, without a dedicated statistical comparison of these diagnostic differences, they should be regarded as mechanistic evidence rather than proof of a unique causal explanation.

4.3. Implications for Three-Dimensional Measurement

The downstream reconstruction results reveal why apparently modest wrapped-phase differences can be important in practice. The aggregate unseen-domain depth MAE decreases from 2.3982 mm for Method 3 to 0.9448 mm for PCIE-FPP (Table 4), a 60.6% reduction. On the representative deeply grooved rack, the wrapped-phase MAEs of Method 3 and PCIE-FPP are relatively close (0.028 and 0.026 rad), whereas their depth MAEs differ markedly (3.23 and 0.16 mm; Figure 7). This disproportionate change is consistent with the sensitivity of multi-frequency fringe-order assignment to localized phase discontinuities: an edge-specific error can lead to an incorrect unwrapping order and a much larger depth residual. The observed suppression of localized error spikes and reflectance-step oscillations (Figure 8c) is therefore potentially more consequential for geometric fidelity than a small improvement in global phase MAE alone.

These benefits are especially relevant to mixed-reflectance parts with sharp transitions, such as the helical gear and grooved rack examined here. Nevertheless, millimeter-valued errors depend on the optical baseline, calibration, object distance, and unwrapping procedure; the particular depth MAEs obtained with this apparatus are not transferable as absolute accuracy guarantees for other FPP systems. Moreover, the reported 28.2 ms inference time pertains to one 1024×1024 fringe image on an NVIDIA RTX 5080. It does not include projector pattern acquisition, three-frequency absolute-phase recovery, or all system-level communication and reconstruction overhead, so it should not be interpreted as a measured end-to-end three-dimensional frame rate.

4.4. Scope, Limitations, and Future Work

The proposed photometric invariance has a defined physical scope. It assumes a fixed calibrated geometry, sufficiently observable fringe modulation, and imaging conditions for which changes in illumination or reflectance predominantly affect the background and modulation envelope. Severe sensor clipping, very weak modulation, cast shadows, multiple reflections, or strongly nonlinear specular responses may violate these assumptions or irreversibly remove useful phase information. Consistent with this distinction, the evaluation excludes pixels with $B_{\mathrm{gt}} \leq 5$ DN or saturated phase-shifting samples (Section 2.6); the reported errors thus characterize phase recovery over the selected observable region rather than recovery of all pixels. Because photometric perturbations can also change the valid-pixel proportion, future studies should report accuracy jointly with coverage when comparing increasingly extreme domains.

The evidence is also bounded by the present evaluation design: 28 unseen test objects, five held-out target photometric domains, and phase-shift diagnostics on eight objects. Although the dual-axis split limits object-identity and target-domain leakage, it does not establish transfer across independently calibrated camera–projector pairs, substantially different fringe frequencies, or previously untested materials. Furthermore, the baseline models differ in architecture and encoding, and aggregate phase/depth comparisons are not accompanied by repeated-training uncertainty intervals or statistical tests. More extensive multi-seed evaluations, cross-device validation, and controlled studies of nonlinear interreflection would strengthen claims about generality and mechanism. Promising future work includes explicit phase-uncertainty estimation, physically informed treatment of low-signal and multipath regions, and end-to-end timing and accuracy assessment in dynamic structured-light measurements. These extensions would test whether the observed robustness can be maintained when the idealized photometric-invariance assumptions no longer hold.

5. Conclusions

This study proposed PCIE-FPP for photometric-domain generalization in single-frame wrapped-phase recovery. By combining a unit-complex-phasor representation, controlled photometric counterfactuals, phase-shift equivariance, and a photometric reconstruction branch, the method achieved a circular phase MAE of 0.0824 rad and a depth MAE of 0.9448 mm on the specified held-out photometric domains and test objects. The improvements over the compared methods were supported by continuous-domain sweeps, physical-response diagnostics, ablation results, and three-dimensional reconstruction experiments. The findings suggest that matching learned phase responses to the physical symmetries of fringe formation can improve robustness with observable photometric changes. Further validation across independent optical systems and more severely nonideal imaging conditions, together with uncertainty and end-to-end timing assessments, remains necessary.

Author Contributions: Conceptualization, M.C.; methodology, M.C.; software, M.C.; validation, M.C., B.H., J.C. and Y.S.; formal analysis, M.C.; investigation, M.C.; resources, M.C.; data curation, M.C.; writing—original draft preparation, M.C.; writing—review and editing, M.C., B.H., J.C. and Y.S.; visualization, M.C. and B.H.; supervision, M.C.; project administration, M.C.; funding acquisition, M.C. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data and code supporting the findings of this study are available from the corresponding author upon request.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Lv, S.; Qian, K. Modeling the measurement precision of fringe projection profilometry. *Light Sci. Appl.* **2023**, *12*, 257. [CrossRef] [PubMed]
2. Lv, S.; Huang, N.; Zou, Y.; Liu, C.; Long, S.Y.; Tang, D.; Jiang, X.; Qian, K. On the precision models of fringe projection profilometry: Unification, simplification and connection. *Light Sci. Appl.* **2026**, *15*, 232. [CrossRef] [PubMed]
3. Dai, R.; Li, W.; Liu, Y.H. Per-pixel calibration based on multi-view 3D reconstruction errors beyond the depth of field. *IEEE Trans. Image Process.* **2025**, *34*, 2124–2132. [CrossRef] [PubMed]
4. Li, Y.; Li, M.; Chen, W.; Qian, J.; Feng, S.; Chen, Q.; Zuo, C. mrPCA-PSP: Robust motion-resistant phase-shifting profilometry based on principal component analysis. *Laser Photonics Rev.* **2025**, *19*, 2401938. [CrossRef]
5. Wan, Y.; Cao, Y.; Xu, M.; Tang, T. Enhanced Fourier–Hilbert-transform suppression for saturation-induced phase error in phase-shifting profilometry. *Opt. Express* **2023**, *31*, 37683–37702. [CrossRef] [PubMed]
6. Han, M.; Chen, W.; Zhang, Q.; Bai, X.; Pan, B. Dual-strategy-guided 2D wavelet transform for single-exposure high dynamic range 3D shape measurement. *Opt. Express* **2025**, *33*, 13145–13160. [CrossRef] [PubMed]
7. Wang, B.; Chen, W.; Qian, J.; Feng, S.; Chen, Q.; Zuo, C. Single-shot super-resolved fringe projection profilometry (SSSR-FPP): 100,000 frames-per second 3D imaging with deep learning. *Light Sci. Appl.* **2025**, *14*, 70. [CrossRef] [PubMed]
8. Chen, W.; Feng, S.; Yin, W.; Li, Y.; Qian, J.; Chen, Q.; Zuo, C. Deep-learning-enabled temporally super-resolved multiplexed fringe projection profilometry: High-speed kHz 3D imaging with low-speed camera. *PhotoniX* **2024**, *5*, 25. [CrossRef]
9. Liu, Y.; Chen, W.; Feng, S.; Xiao, Y.; Jiang, J.; Yu, S.; Liu, Y.; Yin, W.; Chen, Q.; Zuo, C. Temporal super-resolution for high-speed 3D imaging using triple-frequency color-multiplexed fringe projection. *Laser Photonics Rev.* **2026**, *20*, e02920. [CrossRef]
10. Feng, S.; Chen, Q.; Gu, G.; Tao, T.; Zhang, L.; Hu, Y.; Yin, W.; Zuo, C. Fringe pattern analysis using deep learning. *Adv. Photonics* **2019**, *1*, 025001. [CrossRef]
11. Wu, T.; Liu, X.; Zhao, Y.; Qi, L. Single-shot 3-D dense reconstruction using multi-frequency fringe projection for industrial applications. *IEEE Trans. Ind. Inform.* **2024**, *20*, 12747–12757. [CrossRef]
12. Li, J.; Zhang, K.; Luo, L.; Liu, G.; Tang, T.; Wang, Z.; Wan, Y. Deep-learning-enabled single-shot fringe projection profilometry based on inner shifting-phase encoding. *Opt. Express* **2025**, *33*, 49530–49550. [CrossRef] [PubMed]
13. Yin, W.; Che, Y.; Li, X.; Li, M.; Hu, Y.; Feng, S.; Lam, E.Y.; Chen, Q.; Zuo, C. Physics-informed deep learning for fringe pattern analysis. *Opto-Electron. Adv.* **2024**, *7*, 230034. [CrossRef]
14. Li, Z.; Wang, J.; Wang, S.; Zhang, W.; Shan, S.; Yang, Y. Single-shot 3D measurement via deep learning fringe projection profilometry with geometric constraints. *Opt. Laser Technol.* **2025**, *181*, 111735. [CrossRef]
15. Zhang, J.; Luo, B.; Li, F.; Niu, X.; Zhang, Q.; Wang, Y.; Chen, X. Single-exposure optical measurement of highly reflective surfaces via deep sinusoidal prior for complex equipment production. *IEEE Trans. Ind. Inform.* **2023**, *19*, 2039–2048. [CrossRef]
16. Song, X.; Wang, L. Y-FFC Net for 3D reconstruction of highly reflective surfaces. *IEEE Trans. Ind. Inform.* **2024**, *20*, 13966–13974. [CrossRef]
17. Yuan, H.; Li, Y.; Zhao, J.; Zhang, L.; Li, W.; Huang, Y.; Gao, X.; Xie, Q. An adaptive fringe projection method for 3D measurement with high-reflective surfaces. *Opt. Laser Technol.* **2024**, *170*, 110062. [CrossRef]
18. Zhang, M.; Chen, C.; Xie, L.; Zhang, C. Accurate measurement of high-reflective surface based on adaptive fringe projection technique. *Opt. Lasers Eng.* **2024**, *172*, 107820. [CrossRef]
19. Chen, W.; Liu, X.; Ru, C.; Sun, Y. Automated exposures selection for high dynamic range structured-light 3-D scanning. *IEEE Trans. Ind. Electron.* **2023**, *70*, 7428–7437. [CrossRef]
20. Fu, Y.; Fan, J.; Jing, F.; Tan, M. High dynamic range structured light 3-D measurement based on region adaptive fringe brightness. *IEEE Trans. Ind. Electron.* **2024**, *71*, 8080–8090. [CrossRef]
21. Fu, Y.; Fan, J.; Ma, Y.; Jing, F.; Tan, M. An efficient high dynamic range structured light 3-D measurement system based on adaptive brightness control network. *IEEE/ASME Trans. Mechatron.* **2025**, *30*, 2676–2687. [CrossRef]
22. Zhu, Z.; Li, M.; Zhou, F.; You, D. Stable 3D measurement method for high dynamic range surfaces based on fringe projection profilometry. *Opt. Lasers Eng.* **2023**, *166*, 107542. [CrossRef]
23. Xi, D.; Hou, L.; Wu, F.; Qin, Y. Deep learning-based inpainting of high dynamic range fringe pattern for high-speed 3D measurement of industrial metal parts. *Adv. Eng. Inform.* **2024**, *60*, 102428. [CrossRef]
24. Song, K.; Zhang, T.; Sun, C.; Wen, X.; Yan, Y. A novel multi-exposure fusion-induced stripe inpainting method for blade reflection-encoded images. *Adv. Eng. Inform.* **2024**, *60*, 102376. [CrossRef]

25. Lu, L.; Bu, C.; Su, Z.; Guan, B.; Yu, Q.; Pan, W.; Zhang, Q. Generative deep-learning-embedded asynchronous structured light for three-dimensional imaging. *Adv. Photonics* **2024**, *6*, 046004. [CrossRef]
26. Li, X.; Feng, S.; Chen, W.; Jin, Z.; Chen, Q.; Zuo, C. Adaptive structured-light 3D surface imaging with cross-domain learning. *Laser Photonics Rev.* **2025**, *19*, 2401609. [CrossRef]
27. Wang, H.; Zhang, C.; Qian, X.; Wang, X.; Gui, W.; Gao, W.; Liang, X.; Li, X. HDRSL Net for accurate high dynamic range imaging-based structured light 3D reconstruction. *IEEE Trans. Image Process.* **2025**, *34*, 5486–5499. [CrossRef] [PubMed]
28. Shen, M.; He, L.; Zhang, H.; Ma, L.; Li, Y. Deep learning based measurement accuracy improvement of high dynamic range objects in fringe projection profilometry. *Opt. Express* **2024**, *32*, 35689–35702. [CrossRef] [PubMed]
29. Chen, Y.; Yao, P.; Gao, R.; Zhang, W.; Gai, S.; Yu, J.; Da, F. Error model and concise temporal network for indirect illumination in 3D reconstruction. *IEEE Trans. Image Process.* **2024**, *33*, 5849–5863. [CrossRef] [PubMed]
30. Wang, F.; Wang, C.; Guan, Q. Single-shot fringe projection profilometry based on deep learning and computer graphics. *Opt. Express* **2021**, *29*, 8024–8040. [CrossRef] [PubMed]
31. Ren, J.; Tan, C.; Song, W. Hybrid encoding fringe and simulation-to-real scene approach for accurate depth estimation in fringe projection profilometry. *Opt. Express* **2025**, *33*, 14716–14736. [CrossRef] [PubMed]
32. Jiang, H.; Zhu, Z.; Peng, S.; Hu, X.; Peng, Y. Structured light 3D imaging via fusion of spatial-frequency information and domain-adaptive adversarial learning. *Opt. Express* **2025**, *33*, 25378–25396. [CrossRef] [PubMed]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.