PINN-KK: A Physics-Informed Neural Network for Causality-Compliant Spectroscopic Phase Retrieval in Low-Loss Regimes

Olatinwo Adenike Sola[1], Salako Najeem Abiodun[2]

[1]Department of Physics and Engineering Physics, Obafemi Awolowo University, Ile-Ife, Nigeria

[2]Department of Physics and Engineering Physics, Obafemi Awolowo University, Ile-Ife, Nigeria
Corresponding Author Email: *nasalako[at]pg-student.oauife.edu.ng*

Abstract: ***Spectroscopic phase retrieval becomes difficult when weak absorption, finite measurement bandwidth, and noise limit the information available for reconstruction. This study presents PINN-KK, a physics-informed neural-network framework that combines supervised in-bandwidth phase learning with a low-loss, physics-motivated out-of-bandwidth phase constraint and adaptive GradNorm loss weighting. The method is evaluated using synthetic Lorentz-oscillator spectra through an ablation experiment, 28 bandwidth and SNR conditions, and three low-loss thresholds. Removing the physics-based loss produces an approximately four-order-of-magnitude deterioration in out-of bandwidth phase RMS at the reference condition. Across the evaluated conditions, PINN-KK maintains strong in-bandwidth accuracy while substantially reducing out-of-bandwidth phase error relative to Classical KK and the supervised baseline. Performance also improves as the imposed low-loss threshold becomes stricter. These synthetic-data results demonstrate the value of incorporating an explicit low-loss physical prior into neural-network phase retrieval and motivate subsequent validation using experimental spectra.***

Keywords: Kramers-Kronig relations, low-loss optics, phase retrieval, physics-informed neural networks, spectroscopic inversion

1. Introduction

In many spectroscopic applications such as electron energy-loss spectroscopy (EELS) [1], coherent anti-Stokes Raman spectroscopy (CARS) [2,3,4], Fourier transform infrared spectroscopy (FTIR) [5], and terahertz time-domain spectroscopy (THz-TDS) [6,7], phase retrieval in low-loss regimes is ill-posed where weak absorption signals still need to produce reliable dispersion information. These instruments measure only the amplitude (the modulus of the complex dielectric function) but the phase needed for thin-film characterisation, material identification, and dispersion engineering is not directly accessible [8,9]. Amplitude and phase form a Hilbert-transform pair, a relationship that in principle allows the phase to be reconstructed from the amplitude [10,11].

In this work, we present PINN-KK, a physics-informed neural network in which a low-loss approximation of the KK causality condition is embedded as a differentiable physics loss term, enabling phase recovery both inside and outside the measurement window. A GradNorm adaptive weighting scheme automatically balances the physics and data loss terms during training [17,18], enabling reproducible evaluation across 28 BW×SNR conditions and three low-loss threshold values. This paper makes three specific contributions:

1) A low-loss approximation of the KK causality condition, embedded as a differentiable physics loss term enforcing $\varphi \approx 0$ outside the measurement bandwidth;
2) A GradNorm adaptive weighting scheme that automatically balances the physics and data losses during training, eliminating the need for manual λ tuning; and
3) Systematic evaluation across 28 BW×SNR conditions and three low-loss threshold values confirming generalisation and threshold robustness.

The remainder of the paper proceeds as follows. Section II reviews related work; Section III defines the problem; Section IV presents the PINN-KK methodology and experimental design; Section V reports and discusses results across all three stages; Section VI concludes; Section VII outlines future scope.

2. Literature Survey

Some relevant prior work exists in data-driven spectroscopic phase retrieval. Liu et al. [4] demonstrated broadband CARS phase retrieval using a time-domain KK transform, but this approach requires full-bandwidth data and does not generalise to truncated spectra. Klokkou et al. [6] and Zhou et al. [7] applied deep neural networks to THz-TDS refractive index extraction and achieved strong in-bandwidth results, yet neither method imposes a physics constraint that would enforce causality outside the measurement window. A subsequent extension of this approach to two-dimensional materials by Beddoes et al. [16] likewise relies on purely data-driven extraction without a causality constraint. The general PINN framework of Raissi et al. [17] is aimed at spatiotemporal PDE-constrained problems rather than spectral integral inversion, and the adaptive loss balancing scheme GradNorm, introduced by Chen et al. [18], was developed for multi-task learning rather than physics-constrained spectroscopy. PINN-KK addresses these gaps by embedding the KK causality integral as a differentiable loss term with GradNorm-adaptive weighting, enabling phase recovery outside the observed bandwidth without requiring full-spectrum data.

3. Problem Definition

For the past five decades, the standard approach has been the application of classical Kramers-Kronig (KK) inversion, a matrix method that recovers phase from the logarithm of amplitude over the entire frequency axis (i.e. $0 \rightarrow \infty$), whereas real optical instruments only possess finite bandwidth measurements. Classical KK theoretically guarantees a consistent relationship between amplitude and phase via the Cauchy principal value integral. When we truncate such integral, we introduce error of up to ±1.5–2.0 rad at the bandwidth edges [12]. Consequently, this makes the conventional KK inversions violate the causality condition. By causality, we mean a situation where every output must have a corresponding input. Specifically, for a given physical dielectric system $\varepsilon(\omega)$, the causality condition means that such system cannot be polarised before the driving electric field arrives. In other words, this reflects the relativistic constraint that no signal or particle can travel faster than the vacuum speed of light. Also, when a complex refractive index is analytic (holomorphic) in the upper complex plane, its real part can be recovered from the absorption [13]. Likewise, a causality condition can only hold if a corresponding dispersion relation exists [14].

On the other hand, supervised learning provides an alternative under truncated, sparse, and noisy measurement conditions. This is because a pure supervised network learns to interpolate data but it lacks physics built-in, making the predicted phase outside the measurement window unconstrained; such unconstrained predictions may fail to satisfy the imposed physical relation [15].

4. Methodology / Approach

This section describes the complete PINN-KK methodology: the physical model used to simulate ground truth spectra, the mathematical framework of Kramers-Kronig phase retrieval, the network architecture, the physics-informed loss formulation, and how training was carried out. An overview of the full PINN-KK architecture is given in Fig. 1.

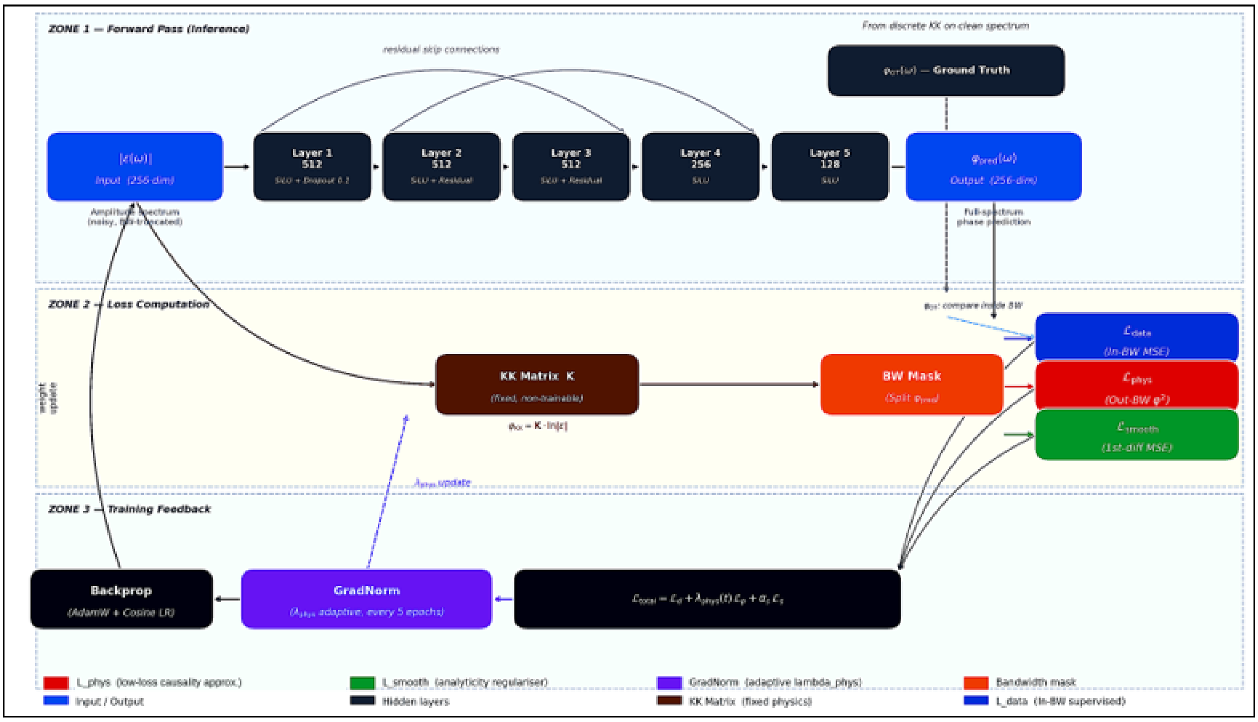

Figure 1: PINN-KK architecture. The 256-dimensional amplitude spectrum $|\varepsilon(\omega)|$ passes through five hidden layers [512, 512, 512, 256, 128] with SiLU activation and residual skip connections to produce the full-spectrum phase $\varphi_{pred}(\omega)$. A bandwidth mask splits φ_{pred} into in-BW and out-BW regions. Three loss terms are computed: L_{data} (in-BW supervised MSE against φ_{GT}), L_{phys} (out-BW low-loss causality-approximation penalty), and L_{smooth} (analyticity regulariser). GradNorm adapts λ_{phys} every 5 epochs. The combined L_{total} is backpropagated via AdamW with cosine learning rate scheduling.

a) The Lorentz Oscillator Model and Ground Truth Simulation

The performance of PINN-KK was evaluated against synthetic spectra generated from the Lorentz Oscillator Model (LOM), which served as the ground truth. The LOM is a well-established physical model for the complex dielectric response of condensed matter systems which is given by

$$\varepsilon(\omega) = \varepsilon_\infty + \Sigma_j \frac{S_j \omega_{0j}{}^2}{\omega_{0j}{}^2 - \omega^2 - i\gamma_j\omega} \tag{1}$$

where $\varepsilon(\omega)$ is the complex dielectric function, ω is the angular frequency, $\varepsilon_\infty = 3.0$ is the high-frequency background dielectric constant, S_j is the oscillator strength, ω_{0j} is the resonance frequency, and γ_j is the damping coefficient. For a single oscillator the real and imaginary parts are:

$$\varepsilon_1(\omega) = \varepsilon_\infty + \frac{S\,\omega_0{}^2(\omega_0{}^2 - \omega^2)}{(\omega_0{}^2 - \omega^2)^2 + (\gamma\omega)^2} \tag{2}$$

$$\varepsilon_2(\omega) = \frac{S\,\omega_0{}^2(\gamma\omega)}{(\omega_0{}^2 - \omega^2)^2 + (\gamma\omega)^2} \tag{3}$$

Three oscillators were used, positioned at normalised resonance frequencies of 1.5, 3.0, and 4.8, with base oscillator strength $S_0 = 0.0193$ and damping $\gamma = 4.88$ (first two) and $\gamma = 5.14$ (third). The frequency grid comprised $N = 256$ equally spaced points over [0.5, 5.5]. Oscillator parameters were randomly jittered per sample, strength scaled over (0.50, 2.00) × S_0, resonance frequencies shifted ±12%, and damping varied over (0.90, 1.50) × γ_0. These parameter ranges were chosen to reproduce the dielectric response of lightly damped ionic crystals in the mid-infrared, for which the low-loss condition is physically satisfied [11].

b) Kramers-Kronig Phase Retrieval

The KK relations arise from the causality condition which requires that a material cannot respond before it is stimulated. For any complex dielectric function that is analytic in the upper half-plane, this causality requirement links its real and imaginary parts through a Hilbert transform. Writing ε(ω) in polar form and taking the natural logarithm gives

$$ln\, \varepsilon(\omega) = ln|\varepsilon(\omega)| + i\varphi(\omega) \tag{4}$$

Applying the Hilbert transform yields:

$$\varphi(\omega) = -\frac{2\omega}{\pi}\mathcal{P}\int_0^\infty \frac{ln|\varepsilon(\omega\prime)|}{\omega\prime^2 - \omega^2}d\omega' \tag{5}$$

$$ln|\varepsilon(\omega')| = -\frac{2}{\pi}\mathcal{P}\int_0^\infty \frac{\omega\prime\, ln|\varepsilon(\omega\prime)|}{\omega\prime^2 - \omega^2}d\omega' \tag{6}$$

where $\mathcal{P}$ denotes the Cauchy principal value. Together, Equations (5) and (6) show that the amplitude over the full frequency range fixes the phase uniquely, and vice versa. The practical limitation that most real instruments deliver only finite bandwidth is the central problem addressed by this work.

c) Discrete KK Matrix Implementation

Equation (4) is discretised into matrix form. The KK matrix K is defined element-wise as:

$$K[i,j] = -\frac{2}{\pi}\frac{\omega_j \Delta\omega}{\omega_j^2 - \omega_i^2}; \quad K[i,i] = 0 \tag{7}$$

Setting K[i,i] = 0 is the standard numerical approach to handling the Cauchy principal value which also removes the singular diagonal term. The phase spectrum is the matrix-vector product $\varphi_{KK} = K \cdot ln|\varepsilon|$. In the low-loss regime where $\varepsilon_2 \ll \varepsilon_1$, the phase satisfies

$$\varphi(\omega) \approx \frac{\varepsilon_2(\omega)}{\varepsilon_1(\omega)} \approx 0 \tag{8}$$

The low-loss condition

$$\frac{|\varepsilon_2(\omega)|}{|\varepsilon_1(\omega)|} < threshold \tag{9}$$

also guarantees |ε| > 0 everywhere, so ln|ε| remains analytic, which is a mathematical prerequisite for Equation (6) to hold.

d) Data Generation Pipeline

Each training sample was generated through a four-step pipeline designed to simulate realistic low-loss spectroscopic measurement conditions.

Step 1- Ground truth generation: The complex dielectric spectrum ε(ω) was computed from (1) for randomly sampled oscillator parameters. The ground truth phase $\varphi_{GT}(\omega)$ was obtained via the discrete KK matrix Equation (7).

Step 2- Low-loss verification: Every generated sample was screened with a hard rejection criterion:

$$\frac{max|\varepsilon_2(\omega)|}{mean|\varepsilon_1(\omega)|} < threshold \tag{10}$$

Failing samples were immediately discarded and regenerated. The maximum safe oscillator strength was bounded using

$$S_{max} = 0.70 \times thresh \times \gamma_{min} \times \varepsilon_{\infty,min} / \omega_{0,max} \tag{11}$$

where $\gamma_{min} = 4.50$.

Step 3- Noise addition: Additive white Gaussian noise was added to the screened samples using

$$\sigma_{noise} = \sqrt{\left(\frac{mean(|\varepsilon|^2)}{10^{SNR/10}}\right)} \tag{12}$$

$$|\varepsilon_{noisy}| = |\varepsilon_{clean}| + \sigma_{noise} \cdot \mathcal{N}(0,1) \tag{13}$$

A flat white noise model was used to prevent overfitting to any single instrument type.

Step 4- Bandwidth truncation: The central bandwidth fraction of the amplitude spectrum was retained, and the remainder was set to zero. For BW = 70% on a 256-point grid, 179 points were observed and 77 were masked at the spectral edges. The model must recover the phase at those 77 masked points using only the physics constraint.

e) Neural Network Architecture

Both the Supervised network and the PINN-KK network shared the same network backbone. The Multi-Layer Perceptron (MLP) takes a 256-dimensional amplitude spectrum |ε(ω)| as input and produces a 256-dimensional phase spectrum φ(ω) as output. Five hidden layers of widths [512, 512, 512, 256, 128] give approximately 500,000 trainable parameters. The Sigmoid Linear Unit (SiLU) activation was used throughout because it is continuously differentiable and better suited than ReLU for smooth spectral regression targets. Residual skip connections between same-width layers were used to prevent vanishing gradients. All weights were initialised using Xavier uniform initialisation.

f) Physics-Informed Loss Function

The PINN-KK network was trained on a three-term composite loss:

$$\mathcal{L}_{total} = \mathcal{L}_{data} + \lambda_{phys}(t) \cdot \mathcal{L}_{phys} + \alpha_{smooth} \cdot \mathcal{L}_{smooth} \tag{14}$$

L_{data} is the supervised MSE between predicted and GT phase inside the observed bandwidth. L_{phys} enforces the low-loss prior φ ≈ 0 outside the measurement window:

$$\mathcal{L}_{phys} = \frac{1}{N_{out}}\Sigma_{\omega \notin BW}\varphi_{pred}{}^2 \tag{15}$$

This is not an arbitrary regulariser. In the low-loss regime $\varphi_{GT}(\omega) \approx 0$ across the full spectrum, so minimising Equation (14) is equivalent to minimising the MSE against the ground truth outside the bandwidth without ever accessing labelled data there. It is worth stating clearly that L_{phys} does not enforce the complete KK integral relation (Equation (6)); it enforces the low-loss approximation φ ≈ 0 that follows once the accepted samples satisfy the threshold in Equation (9). This approximation is justified here because every training and evaluation sample was screened against the low-loss condition (Equation (10)), which restricts the dataset to materials for which $\varphi_{GT}(\omega)$ is close to, but not necessarily exactly, zero across the full spectrum. L_{smooth} penalises first differences over the full spectrum, enforcing the analyticity requirement that $\varphi(\omega)$ varies smoothly. The Supervised network was trained on L_{data} only while Classical KK applied Equation (7) analytically without training, producing systematic errors at bandwidth edges due to spectral truncation.

g) Gradient Normalisation (GradNorm) Adaptive Weighting

The physics loss weight λ_{phys} in Equation (14) was adapted automatically every 5 epochs using the GradNorm scheme [18]. GradNorm computes the gradient norm ratio

$$\lambda_{target}(t) = \frac{\|\nabla_\theta \mathcal{L}_{data}\|}{\|\nabla_\theta \mathcal{L}_{phys}\|} \tag{16}$$

It updates the weight in log-space via an exponential moving average (EMA) with coefficient 0.90, guaranteeing $\lambda_{phys} > 0$ at all times. In practice, λ_{phys} grows from 0.003 to approximately 0.80 by epoch 80, a ~200× increase reflecting the natural gradient imbalance between L_{data} (rapidly decreasing) and L_{phys} (more slowly decreasing). The two bounds $\lambda_{min} = 10^{-6}$ and $\lambda_{max} = 1.0$ prevent degenerate solutions. The smoothness weight $\alpha_{smooth} = 0.001$ was fixed throughout the training process.

This eliminates manual λ tuning and represents a concrete methodological advance over fixed-λ PINNs [17] where λ must be re-tuned for every new operating condition. Results were reproducible across all 28 BW×SNR conditions without re-tuning λ_{phys}.

h) Optimiser and Training Protocol

Both models were trained using Adaptive Moment Estimation with Weight Decay (AdamW). The learning rate lr = 1×10^{-3} and weight_decay = 1×10^{-5} were used with CosineAnnealingWarmRestarts ($T_0 = 50$, $T_{mult} = 2$, $\eta_{min} = 10^{-6}$). Training ran for a maximum of 350 epochs, with early stopping triggered after 40 consecutive epochs without validation improvement. Each model was trained on $N_{train} = 5{,}000$ synthetic spectra with batch size 64. $N_{val} = 500$ spectra monitored validation loss, and $N_{test} = 1{,}000$ completely held-out spectra were used for all reported evaluations. Each training was subjected to three seeds [42, 137, 314] per condition and the results are reported as mean ± 95% confidence interval (CI) using the t-distribution (df = 2, $t_{crit} = 4.303$). Three seeds represent the minimum reproducible set given the computational cost of 168 training runs; confidence intervals at near-zero OutBW_PhysRMS values are consequently wide relative to the mean and should be interpreted as order-of-magnitude estimates at those conditions.

i) Experimental Design

The experimental evaluation was structured into three progressive stages. Stage 1 isolates the contribution of each loss component through a controlled ablation at a single operating point. Stage 2 tests generalisation across the full range of bandwidth and noise conditions while Stage 3 examines PINN-KK sensitivity to the low-loss threshold parameter. All stages used synthetic spectra with the GT phase known exactly.

Stage 1- Ablation Study

Stage 1 is a controlled ablation study at single operating point BW = 70%, SNR = 15 dB, low-loss threshold = 1.0%, and three seeds. Three variant runs were conducted as described below.

Table I: Ablation Study Run Configurations

Run	L_d	L_p	L_s	Purpose
RunA	✓	✓	✓	Full PINN-KK. All three loss terms active. All paper figures derived from RunA.
RunB	✓	✓	–	Smoothness loss removed. Isolates the independent contribution of L_{smooth}.
RunC	✓	–	–	Physics and smoothness disabled. Negative control: confirms physics drives performance, not architecture.

RunC is our negative control experiment. It uses the identical architecture as RunA but removes all physics constraints, reducing the PINN-KK network to a supervised-only network. Any performance gap between RunA and RunC isolates the contribution of the KK physics loss independent of network capacity.

Stage 2- Bandwidth × SNR Sweep

Stage 2 is our main experimental evaluation. PINN-KK was tested across four bandwidth fractions (30%, 50%, 70%, 100%) and seven SNR levels (10, 15, 20, 25, 30, 35, 40 dB), giving 28 conditions. The low-loss threshold was fixed at 1.0% throughout. For supervised and PINN-KK models each condition ran at three seeds, giving a total of 168 training runs, completed in approximately 10 hours on an HP Core i7 CPU. The BW = 100% conditions serve as sanity-check baselines where OutBW_PhysRMS = 0 by definition for all models.

Stage 3- Threshold Sensitivity Analysis

Stage 3 investigated sensitivity to the low-loss threshold at the fixed operating point BW = 70%, SNR = 15 dB, and three seeds, giving 9 additional runs across the three threshold values described below.

Table II: Stage 3 Threshold Sensitivity Configurations

Threshold	Filter Condition	Physical Interpretation
1.0% (0.010)	$\max\lvert\varepsilon_2\rvert / \text{mean}\lvert\varepsilon_1\rvert < 0.010$	Baseline: weakest low-loss requirement. Broadest material distribution.
0.7% (0.007)	$\max\lvert\varepsilon_2\rvert / \text{mean}\lvert\varepsilon_1\rvert < 0.007$	Intermediate: tighter constraint on material lossiness.
0.5% (0.005)	$\max\lvert\varepsilon_2\rvert / \text{mean}\lvert\varepsilon_1\rvert < 0.005$	Strictest: physics prior strongest; phase closest to zero everywhere.

All other training parameters were held constant from Stage 2. Stage 3 constitutes a new experimental finding; the performance of PINN-KK as a function of material class strictness had not previously been systematically evaluated.

Paper ID: SR26810183613 DOI: https://dx.doi.org/10.21275/SR26810183613

5. Evaluation Metrics

Three metrics are reported. The Out-of-Bandwidth Physics RMS (OutBW_PhysRMS) is the primary metric representing the RMS of predicted phase outside the observed bandwidth, where data are sparsely sampled or not available, directly quantifying physics-law compliance. The Root Mean Square Error Full (RMSE_Full) is computed as the square root of the mean squared deviation between the predicted and ground-truth phase, evaluated over the complete 256-point spectrum. The In-BW RMSE measures accuracy inside the observed bandwidth. All metrics are reported as mean ± 95% CI across three seeds. Improvement ratios are computed as the ratio of OutBW_PhysRMS for Classical KK to that of PINN-KK.

6. Results & Discussion

Three methods have been used for spectroscopic phase retrieval: Classical KK inversion, pure Supervised model, and the PINN-KK model. We benchmark the performance of each model against the GT phase across all evaluated conditions. The primary metric is OutBW_PhysRMS, which measures causality compliance outside the measurement window. We also report RMSE_Full over the complete spectrum and In-BW RMSE to confirm spectral accuracy.

1) Stage 1- Ablation Results

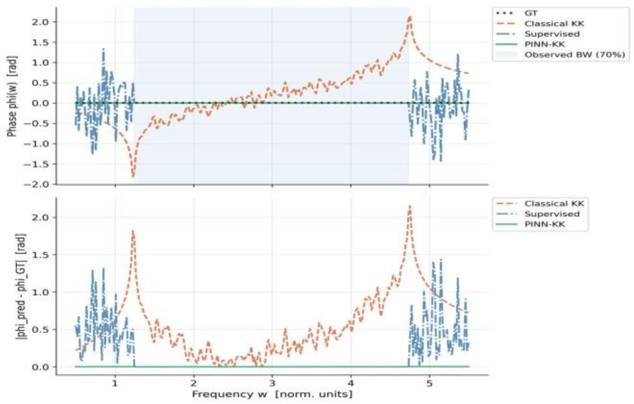

Figure 2: Phase reconstruction at the reference condition (BW = 70%, SNR = 15 dB). Upper panel: predicted phase φ(ω) for PINN-KK (green solid), Classical KK (yellow dashed), and Supervised (blue dashed) against GT (black dashed), with the observed bandwidth shaded. Lower panel: absolute phase error $|\varphi_{pred} - \varphi_{GT}|$ across the full frequency axis.

Fig. 2 represents a plot of OutBW_PhysRMS for 70% bandwidth at 15 dB SNR, which is our reference condition. There are two panels in Fig. 2: the upper and the lower panels. In the upper panel, the dashed black line represents the GT phase, which remains close to zero both inside and outside the measurement window. This GT phase is the discretised KK-matrix solution (Equation (7)) applied to the full-spectrum amplitude.

The PINN-KK model is represented by the green solid line. It is not visible in the upper panel because it sits exactly on the GT line. But in the lower panel, it lies close to zero across the entire frequency axis, confirming that the network's low-loss causality approximation closely matches the GT phase at all frequencies.

The yellow dashed line represents the Classical KK inversion. It oscillates between ±1.8 radians (rad) due to noise and truncation artefacts at the bandwidth edges. The blue dashed line represents the pure Supervised model, which performs well inside the measurement window but generalises arbitrarily between ±1.2 rad outside the measurement window because its training loss lacks an explicit physics constraint. At this reference condition, PINN-KK achieves 88,353× lower OutBW_PhysRMS than Classical KK and 50,923× lower than the supervised model.

Table III shows the summary of our ablation results. This reveals what each loss term in Equation (14) contributes to the training process of PINN-KK before sweeping across different BW and SNR levels. RunA represents full PINN-KK consisting of all three loss terms (L_{data}, L_{phys}, and L_{smooth}). In RunB, we removed the smoothness loss term, and in RunC, which was our negative control experiment, we disabled both the physics and smoothness loss terms. RunC isolates the effect and contribution of the enforced physics in the training process of PINN-KK. Practically, RunC is a similar version of the pure supervised network.

From Table III, the In-BW RMSE for both PINN-KK (RunA, RunB, and RunC) and the pure supervised model is 0.00127, which is better than Classical KK at 0.6262. This is expected because both PINN-KK and the supervised model have access to the same data and are trained under the same conditions. The performance of Classical KK is largely affected by truncation and noise. This also confirms that the physics loss term (L_{phys}) adds zero accuracy cost inside the measurement window.

Table III: Ablation Study Results: BW = 70%, SNR = 15 dB, Thresh = 1.0%

Run	In-BW RMSE	OutBW_PhysRMS
RunA (Full PINN-KK)	$0.00127 \pm 3.5 \times 10^{-5}$	$1.03 \times 10^{-5} \pm 3.9 \times 10^{-6}$
RunB (No L_s)	$0.00127 \pm 3.5 \times 10^{-5}$	1.10×10^{-5}
RunC (No L_p, No L_s)	$0.00127 \pm 3.5 \times 10^{-5}$	0.492
Supervised MLP	0.00127	0.523
Classical KK	0.6262	0.907

The OutBW_PhysRMS for RunA, RunB, and RunC are 1.03×10^{-5}, 1.10×10^{-5}, and 0.492, respectively. It is evident that from RunA and RunC values, there is a four-order-of-magnitude degradation. The OutBW_PhysRMS for the Supervised model is 0.523 (our reference baseline), while Classical KK is 0.907 (our analytical baseline). The OutBW_PhysRMS values of RunC (0.492) and the supervised model (0.523) are approximately equal, confirming RunC as a valid negative control. PINN-KK achieves accuracy 88,353× better than Classical KK and outperforms the Supervised model by 50,923×.

2) Stage 2- Bandwidth × SNR Sweep

PINN-KK and pure Supervised models were evaluated across all combinations of four BW fractions (30%, 50%, 70%, 100%), seven SNR levels (10, 15, 20, 25, 30, 35, 40 dB), and 3 seeds, giving 28×3×2 = 168 total training runs across both models. Though some of the trainings converged before reaching the maximum epoch of 350. The low-loss threshold was fixed at 1.0% throughout. At 30% BW and 20 dB SNR, PINN-KK showed an improvement of 7,804× better than Classical KK, while at 50% BW and 40 dB SNR the improvement increases to 238,422×. The Stage 1 reference condition was exactly replicated in Stage 2 at BW = 70% and SNR = 15 dB yielding RMSE_Full = 0.001895, compared to Stage 1 which is 0.001897, a difference of only 2×10^{-6}, confirming reproducibility across stages. This performance gap is consistent with the in-bandwidth-only training objective of the neural network methods reported by Klokkou et al. [6] and Zhou et al. [7], which impose no out-of-bandwidth constraint and would therefore yield OutBW_PhysRMS values comparable to the Supervised model baseline reported here.

3) C. SNR Robustness

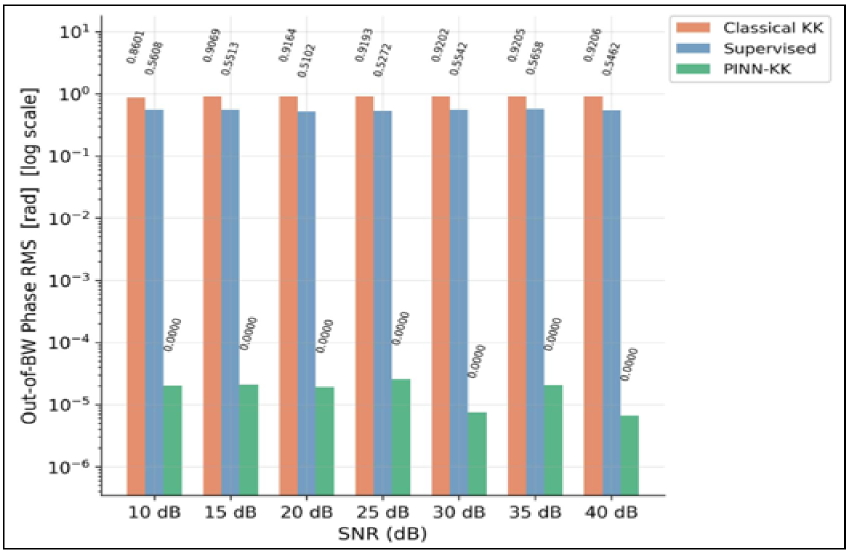

Figure 3: OutBW_PhysRMS on a log scale versus SNR (10–40 dB) at BW = 70% for all three models. PINN-KK (green) remains near 10^{-5} across all SNR values while Classical KK (orange) and Supervised (blue) sit near 0.9 and 0.55 respectively.

Fig. 3 shows the log scale of OutBW_PhysRMS across various values of SNR = 10–40 dB (noisy to most clean data) at 70% for all three models. At each SNR value, the PINN-KK RMSE is of the order of 10^{-5}, with Classical KK and the pure Supervised model sitting near 0.9 and 0.55 respectively. This shows that PINN-KK is more robust to noisy data than both Classical KK and the pure Supervised model. In addition, PINN-KK RMSE_Full variance within each BW row is less than 4×10^{-6} across all 7 SNR values. The Supervised model fluctuates between 0.23–0.42. This confirms PINN-KK is essentially SNR-independent.

4) Bandwidth Degradation

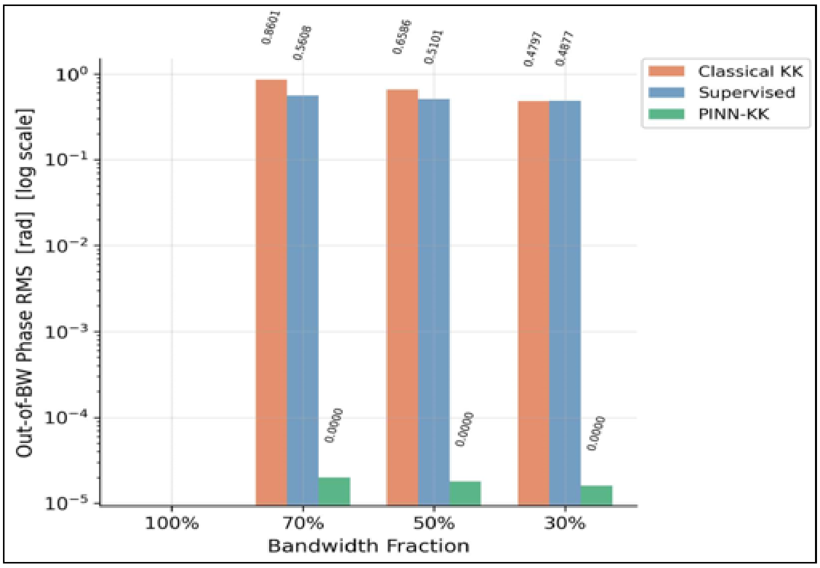

Figure 4: OutBW_PhysRMS on a log scale versus bandwidth fraction (30%, 50%, 70%) at SNR = 10 dB (the most challenging noise condition) for all three models. PINN-KK (green) remains near 1.5×10^{-5} across all BW fractions. BW = 100% excluded (OutBW_PhysRMS = 0 by definition for all models at full bandwidth).

Paper ID: SR26810183613 DOI: https://dx.doi.org/10.21275/SR26810183613

Fig. 4 shows the log-scale OutBW_PhysRMS against different BW fractions of 100%, 70%, 50%, and 30% for all three models. To check the robustness of PINN-KK against BW degradation, we deliberately use the most challenging noisy condition (10 dB) to show how PINN-KK model dominates over Classical KK and the pure Supervised model. From this figure, the value of PINN-KK model across all the bandwidth fractions is approximately 1.5×10^{-5}. At BW = 70%, Classical KK is 0.86 (PINN-KK shows an improvement of approximately 57,000×) versus Supervised at 0.56. At BW = 50%, Classical KK is 0.66 versus Supervised at 0.51. While at BW = 30%, Classical KK is 0.48 versus Supervised at 0.49. This confirms that the PINN-KK model does not degrade as bandwidth narrows. Classical KK degrades from 0.48 to 0.86 as BW increases. Supervised shows a degradation from 0.49 to 0.56 as BW fraction increases.

5) Anomalous Condition- BW = 30%, SNR = 20 dB

Our experiment witnessed an anomalous condition at BW = 30%, SNR = 20 dB. At this condition, PINN-KK model exhibited the highest variance across seeds (CV, the coefficient of variation, σ/μ across three seeds, is 1.59). The mean value of the OutBW_PhysRMS nevertheless remains 7,804× below the Classical KK baseline.

6) GradNorm Training Dynamics

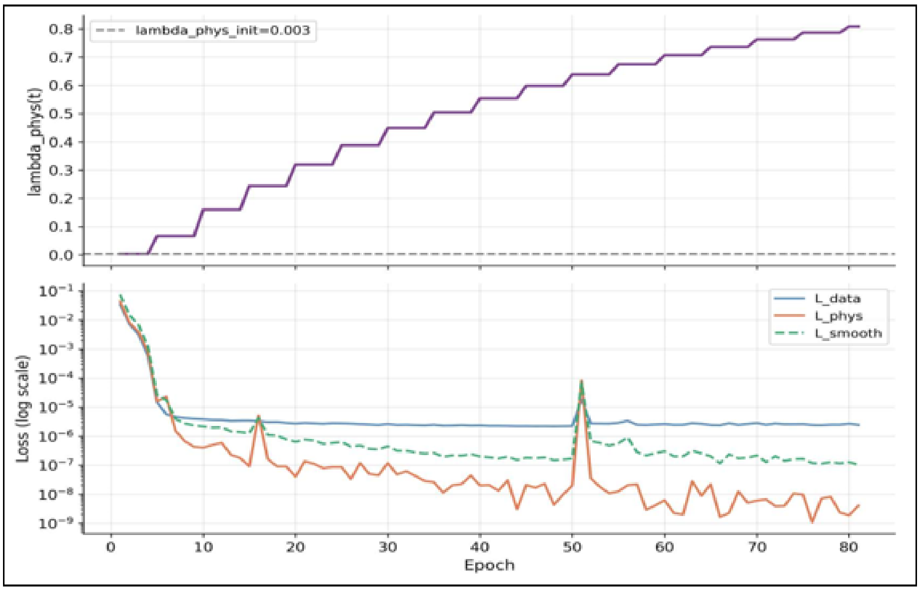

Figure 5: GradNorm training dynamics at BW = 70%, SNR = 15 dB. Upper panel: λ_{phys} weight evolution from 0.003 to ~0.80 across epochs. Lower panel: log-scale loss curves for L_{data}, L_{phys}, and L_{smooth}. The cosine warm-restart at epoch 50 produces a visible spike before final convergence

Fig. 5 shows two different panels. The upper panel shows how the λ_{phys} weight grows from 0.003 to approximately 0.80 by epoch 80. The lower panel shows the log scale plot of loss versus epoch for L_{data}, L_{phys}, and L_{smooth}. At epoch 50, the cosine warm-restart produces a clear spike in both panels for all the loss terms before final convergence. Early stopping fires at approximately 90 epochs for most conditions, even before reaching the 350-epoch maximum cap. The plot reveals how three loss terms converge consistently, with L_{phys} reaching 10^{-8} to 10^{-9} at convergence. This is direct evidence that GradNorm correctly up-weights the physics constraint as data loss decreases faster than physics loss.

7) Stage 3- Threshold Sensitivity

Paper ID: SR26810183613 DOI: https://dx.doi.org/10.21275/SR26810183613

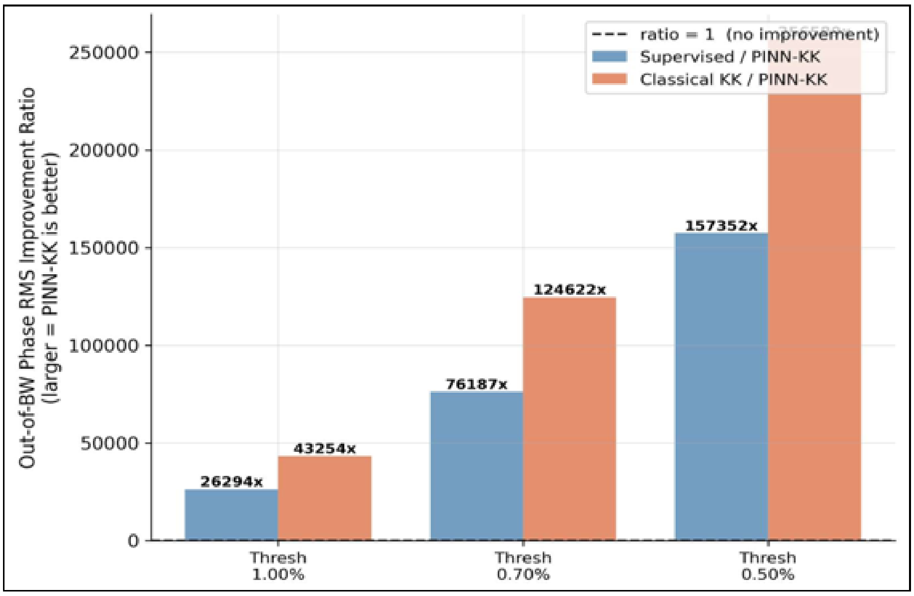

Figure 6: Improvement ratios of PINN-KK over Classical KK (orange bars) and over the Supervised model (blue bars) at each low-loss threshold. Both ratios grow monotonically: Classical KK ratio 43,254× → 124,622× → 256,750×; Supervised ratio 26,294× → 76,187× → 157,352× as threshold tightens from 1.0% to 0.5%.

Fig. 6 shows how the PINN-KK improvement ratio against each baseline grows as the low-loss threshold tightens. Against Classical KK, the ratio increases from 43,254× at thresh = 1.0% to 124,622× at 0.70%, and reaches 256,750× at 0.5%. Against the Supervised model, the corresponding ratios are 26,294×, 76,187×, and 157,352×. In both cases the improvement is strictly monotonic, confirming that a stricter low-loss threshold simultaneously improves PINN-KK's absolute performance and widens its margin over both baselines.

The absolute PINN-KK OutBW_PhysRMS values underline these ratios: $2.097 \times 10^{-5} \rightarrow 7.27 \times 10^{-6} \rightarrow 3.53 \times 10^{-6}$ as reported in Table IV.

Table IV:
Threshold Improvement Ratios for PINN-KK: BW = 70%, SNR = 15 dB

Threshold	OutBW_ PhysRMS	vs Classical KK	RMSE_ Full	In-BW RMSE
1.0% (0.010)	2.10×10^{-5}	43,254×	0.001895	0.001274
0.7% (0.007)	7.27×10^{-6}	124,622×	0.00133	0.000893
0.5% (0.005)	3.53×10^{-6}	256,750×	0.000949	0.000638

Table IV shows the improvement ratio of PINN-KK model over Classical KK outside the observed BW, inside the observed BW, and full spectrum (RMSE_Full). As the threshold tightens from 1% to 0.5%, the RMSE_Full halves exactly from 0.001895 to 0.000949, a 2.00× improvement. The In-BW RMSE is identical to the supervised model at every threshold value, confirming zero accuracy cost across all threshold values. Note that Stage 3 thresh = 1.0% yields OutBW_PhysRMS = 2.097×10^{-5}, which differs from Stage 1 RunA (1.03×10^{-5}) at the same nominal condition; this difference arises from independent random seeds and a fresh data draw and is statistically consistent, not contradictory.

7. Discussion

1) Physics Loss Is the Critical Innovation

The ablation results establish that L_{phys} is the critical component of PINN-KK. The four-order-of-magnitude collapse in RunC confirms that L_{phys} drives the out-of-bandwidth result but not the network architecture. The identical In-BW RMSE across RunA, RunB, RunC, and the supervised model proves that the physics constraint adds no accuracy cost inside the measurement window. The network architecture is necessary but not sufficient; the physics constraint is precisely what produces causality compliance outside the bandwidth.

Comparing RunA (1.03×10^{-5}) against RunC (0.492) reveals a factor of 47,754× degradation attributable solely to the removal of L_{phys}. The OutBW_PhysRMS values of RunC (0.492) and the supervised model (0.523) are approximately equal, confirming RunC as a valid negative control. The architecture alone, without physics supervision, performs no better than a purely supervised network.

2) GradNorm Enables Stable Multi-Loss Training

The λ_{phys} weight growth from 0.003 to approximately 0.80 reflects the natural gradient imbalance that develops during training. L_{data} decreases rapidly in early epochs as the network fits the in-bandwidth phase, while L_{phys} decreases more slowly because the out-of-bandwidth constraint is harder to enforce. GradNorm [18] automatically compensates by increasing λ_{phys} to maintain comparable gradient magnitudes from both terms, preventing the physics constraint from being overwhelmed by the data loss. This is direct evidence that GradNorm correctly up-weights the physics constraint as data loss decreases faster than physics loss.

The adaptive mechanism eliminates the need for manual λ tuning and constitutes a concrete methodological advance over fixed-λ PINN formulations [17], where λ must be re-tuned for every new operating condition. The GradNorm approach produced consistent, reproducible results across all

Paper ID: SR26810183613 DOI: https://dx.doi.org/10.21275/SR26810183613

84 PINN-KK training runs (28 conditions × 3 seeds) in Stage 2 without any manual adjustment of λ_{phys}.

3) Generalisation Across BW, SNR, and Threshold

Stage 2 confirms that PINN-KK holds its performance advantage across all 28 BW×SNR conditions, demonstrating that the model does not overfit. The SNR robustness result (RMSE_Full variance less than 4×10^{-6} across all noise levels) suggests that the physics constraint acts as an implicit noise regulariser by forcing the out-of-bandwidth phase toward zero regardless of the noise realisation. L_{phys} stabilises the extrapolated phase against noise-induced fluctuations that would otherwise corrupt the supervised extrapolation.

The irregularity at BW = 30%, SNR = 20 dB is a consequence of and consistent with the compounding of maximum bandwidth truncation and transition-zone noise where GradNorm weight adaptation is most sensitive. However, the mean OutBW_PhysRMS remains 7,804× below the Classical KK baseline, confirming that even under the most adverse conditions the physics constraint retains decisive practical advantage.

Stage 3 reveals a previously unknown property where stricter low-loss threshold improves PINN-KK performance monotonically. A stricter threshold means the training dataset contains materials whose phase is even closer to zero everywhere, so the physics prior ($\varphi \approx 0$ outside bandwidth) is more precisely satisfied. This is further confirmed by CV dropping to 0.12 at thresh = 0.5%.

8. Limitations

The following limitations apply to the current implementation of PINN-KK. First, we used synthetic Lorentz data to simulate the performance of PINN-KK and generalisation to real experimental spectra such as FTIR and THz have not been tested. This is the subject of ongoing work in Stage 4. Second, the low-loss regime restriction is not an arbitrary limitation but a mathematical requirement for $\ln|\varepsilon|$ to remain analytic. It restricts applicability to low-loss materials and excludes strongly absorbing media where $|\varepsilon_2|$ is comparable to $|\varepsilon_1|$ across the spectrum. Third, we used N = 3 seeds, which produced wide confidence intervals for near-zero OutBW_PhysRMS values. We plan to use $N \geq 5$ seeds in future work to tighten statistical confidence intervals, particularly at the strictest threshold values where PINN-KK performance is strongest.

9. Conclusion

This study demonstrates a physics-informed neural-network approach for phase retrieval from synthetic low-loss Lorentz spectra under finite-bandwidth and noisy conditions. The ablation experiment shows that the physics-based loss is responsible for the substantial reduction in out-of-bandwidth phase error, while in-bandwidth accuracy remains comparable to the supervised baseline. Evaluation across bandwidth and SNR conditions indicates stable performance within the investigated synthetic domain, and the threshold analysis shows further improvement as the low-loss criterion becomes stricter. GradNorm also provides adaptive balancing of the data and physics-based objectives without condition-specific manual adjustment of the physics-loss weight. These findings support PINN-KK as a promising approach for low-loss phase retrieval, while validation on experimentally measured spectra remains necessary before broader practical conclusions can be established.

10. Future Scope

Finally, beyond evaluation on synthetic data, one primary open question remains the generalisation of PINN-KK to real experimental spectra, where noise characteristics and oscillator distributions differ from the synthetic Lorentz model used here. Our future work will be to validate the performance and accuracy of PINN-KK on real FTIR or THz experimentally measured data.

Acknowledgements
The authors acknowledge the Department of Physics and Engineering Physics, Obafemi Awolowo University, Ile-Ife, for providing the academic and computational environment in which this work was carried out. All model training was performed on locally available computing resources (HP Core i7 CPU), without external funding.

Competing Interests
The authors declare that no competing interests exist.

Authors' Contributions
This work was carried out in collaboration between both authors. Author OAS supervised the research, provided guidance on the theoretical framing of the Kramers-Kronig causality problem, and reviewed and revised the manuscript. Author SNA conceived the PINN-KK methodology, implemented the network architecture and physics-informed loss functions, designed and executed the Stage 1–3 experiments, performed the data analysis, generated all figures and tables, and drafted the manuscript. Both authors read and approved the final manuscript.

Consent
Not applicable. This study did not involve human participants, and no personal or identifiable data of any individual were used.

Ethical Approval
Not applicable. This study did not involve human participants, human data, or animals. All spectra used in this work were synthetically generated from the Lorentz oscillator model described in Section IV-A, and no ethical clearance was therefore required.

References

[1] Egerton RF. Electron Energy-Loss Spectroscopy in the Electron Microscope. 3rd ed. New York: Springer; 2011.
[2] Houhou R, Rosendahl P, Pajer P, Wiest A, Uckermann O, Schackert G, Galli R, Meyer T. Deep learning-based classification of cancer cells using coherent anti-Stokes Raman scattering (CARS) microscopy. Sci Rep. 2020; 10: 12904.
[3] Chen BC, Lim SH. Optimal laser pulse shaping for interferometric multiplex coherent anti-Stokes Raman

scattering microscopy. J Phys Chem B. 2008;112(12):3653–3661.

[4] Liu Y, Lee YJ, Cicerone MT. Broadband CARS spectral phase retrieval using a time-domain Kramers-Kronig transform. Opt Lett. 2009;34(9):1363–1365.

[5] Griffiths PR, de Haseth JA. Fourier Transform Infrared Spectrometry. 2nd ed. Hoboken: John Wiley & Sons; 2007.

[6] Klokkou N, Gorecki J, Beddoes B, Apostolopoulos V. Deep neural network ensembles for THz-TDS refractive index extraction exhibiting resilience to experimental and analytical errors. Opt Express. 2023;31(26):44575–44588.

[7] Zhou Z, Jia S, Cao L. A general neural network model for complex refractive index extraction of low-loss materials in the transmission-mode THz-TDS. Sensors. 2022;22(20):7877.

[8] Tyson RK, Frazier BW. Principles of Adaptive Optics. 5th ed. Boca Raton: CRC Press; 2022.

[9] Colomb T, Kühn J. Digital holographic microscopy. In: Leach R, editor. Optical Measurement of Surface Topography. Berlin: Springer; 2011. p. 209–235.

[10] Whittaker KA, Keaveney J, Hughes IG, Adams CS. Hilbert transform: applications to atomic spectra. Phys Rev A. 2015;91(3):032513.

[11] Lucarini V, Saarinen JJ, Peiponen KE, Vartiainen EM. Kramers-Kronig Relations in Optical Materials Research. Berlin: Springer; 2005.

[12] Harbecke B. Application of Fourier's allied integrals to the Kramers-Kronig transformation of reflectance data. Appl Phys A. 1986; 40: 151–158.

[13] Kramers HA. La diffusion de la lumière par les atomes. Atti Congr Int Fis Como. 1927; 2: 545–557.

[14] Kronig R de L. On the theory of dispersion of X-rays. J Opt Soc Am. 1926; 12:547–557.

[15] Gralak B, Lequime M, Zerrad M, Amra C. Phase retrieval of reflection and transmission coefficients from Kramers-Kronig relations. J Opt Soc Am A. 2015;32(3):456–462.

[16] Beddoes B, Klokkou N, Gorecki J, Whelan PR, Bøggild P, Jepsen PU, Apostolopoulos V. THz-TDS: extracting complex conductivity of two-dimensional materials via neural networks trained on synthetic and experimental data. Opt Express. 2025;33(7):14872–14884.

[17] Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. J Comput Phys. 2019; 378: 686–707.

[18] Chen Z, Badrinarayanan V, Lee CY, Rabinovich A. GradNorm: gradient normalization for adaptive loss balancing in deep multitask networks. In: Proceedings of the 35th International Conference on Machine Learning; 2018. p. 794–803.