# APEX: AMPLITUDE ANCHORS AND PHASE PRIORS FOR TARGET-SCARCE HIGHER-FREQUENCY WAVE PREDICTION

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Learning-based surrogates have become increasingly effective for wave-field prediction, and neural operators in particular have shown strong performance within observed frequency regimes. However, higher-frequency prediction under scarce target supervision remains comparatively underexplored, especially in wave problems where higher-frequency data are substantially more expensive to simulate or measure than lower-frequency data. A central difficulty is that cross-frequency transfer is inherently asymmetric: coarse amplitude structure remains relatively stable across frequencies, whereas phase-sensitive oscillatory structure deteriorates much more rapidly as frequency increases. Motivated by this asymmetry, we propose **APEX**, **A**mplitude-anchored and **P**hase-prior-guided **E**nhancement from e**X**trapolated coarse predictions, a framework for target-scarce higher-frequency wave-field prediction. A lower-frequency neural operator first provides a coarse prediction in the target-frequency regime, from which we retain only the amplitude as a transferable structural anchor. A conditional flow-matching enhancer then reconstructs the target higher-frequency field under the guidance of a Green's-function-inspired phase prior. Experiments on SimpleWave, Helmholtz, and Maxwell benchmarks show that APEX consistently outperforms direct lower-to-higher extrapolation, target-adapted operator, and joint generative baselines under limited target-frequency supervision. Our results suggest that reliable higher-frequency prediction of oscillatory wave fields should not rely on direct end-to-end transfer of the full complex field, but instead on explicitly reusing transferable coarse structure while separately recovering the missing oscillatory detail.

## 1 INTRODUCTION

Learning-based surrogates have become a powerful paradigm for approximating solution maps of parameterized PDEs, with broad impact across fluid dynamics, electromagnetics, acoustics, and other scientific domains Gu et al. (2022); Zou et al. (2024); Hao et al. (2023). Among these approaches, neural operators have emerged as a particularly effective class of models for learning nonlocal mappings in function space, making them a natural backbone for wave-field prediction Li et al. (2021); Kovachki et al. (2023); Lu et al. (2021). More broadly, this places wave-field modeling in the growing problem of *cross-frequency generalization* for learned surrogates: models are expected not only to interpolate within observed spectral regimes, but also to remain reliable as the queried frequency changes. In many practical wave problems, this challenge is especially important at higher frequencies, which are much more expensive to simulate or measure Ihlenburg & Babuška (1995).

Despite rapid progress in learned PDE surrogates and neural operators, this higher-frequency regime remains comparatively underexplored. Most existing studies on cross-frequency wave prediction focus on interpolation within observed spectral ranges, adaptation with target-regime supervision, or general in-range surrogate accuracy Gu et al. (2022); Seo et al. (2024); Sun et al. (2025); Wang et al. (2024). In this paper, we focus on a particularly important regime of this broader problem: *higher-frequency prediction under scarce target supervision*. This setting arises naturally in wave problems because higher-frequency samples typically require finer discretization, stricter numerical control, and substantially higher simulation or measurement cost Ihlenburg & Babuška (1995). As a



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

result, practitioners often have abundant lower-frequency data but only limited supervision in the higher-frequency target regime.

A straightforward strategy in this setting is to train a surrogate on abundant lower-frequency data and directly apply it at unseen higher frequencies, i.e., *direct lower-to-higher extrapolation*. Another common strategy is to adapt the lower-frequency model using the limited target-frequency data that can be afforded Sun et al. (2025); Wang et al. (2024). In practice, however, both approaches often remain unreliable in highly oscillatory higher-frequency regimes, where wave responses are typically more sensitive to propagation mismatch Keller (1962); Sethian (1996). More importantly, what remains insufficiently understood is not merely whether lower-to-higher prediction degrades, but *which* components of the complex field remain transferable and *which* ones fail first under frequency shift. This missing mechanism-level understanding makes it difficult to design models that generalize reliably beyond the observed regime.

In this work, we study higher-frequency prediction under scarce target supervision through the lens of transferability across field components. Our central finding is that cross-frequency transfer is inherently asymmetric: coarse amplitude structure remains substantially more stable across frequencies, while phase-sensitive oscillatory structure deteriorates much more rapidly as the frequency gap increases. This asymmetry appears both in the ground-truth cross-frequency similarity structure and in the outputs of neural operators trained on lower frequencies and directly queried at higher frequencies. Motivated by this finding, we propose **APEX**, **A**mplitude-anchored and **P**hase-prior-guided **E**nhancement from e**X**trapolated coarse predictions, a framework for target-scarce higher-frequency wave-field prediction. APEX uses a frozen lower-frequency Fourier neural operator (FNO) to provide a coarse amplitude anchor, and then reconstructs the target higher-frequency field using a conditional flow-matching enhancer Lipman et al. (2023); Tong et al. (2024) guided by a Green's-function-inspired phase prior. Experiments on SimpleWave, Helmholtz, and Maxwell benchmarks show that APEX consistently outperforms direct lower-to-higher extrapolation, target-adapted operator, and joint generative baselines under limited target-frequency supervision Wang et al. (2024); Sun et al. (2025); Huang et al. (2024). More broadly, our results suggest that reliable higher-frequency prediction should not rely on direct end-to-end transfer of the full complex field, but instead on reusing transferable coarse structure while separately recovering the missing oscillatory detail.

Our contributions are summarized as follows:

- We identify a key *transfer asymmetry* in higher-frequency wave field prediction: coarse amplitude remains more stable across frequencies than phase-sensitive oscillatory structure.
- We propose **APEX**, a framework that decomposes higher-frequency prediction under scarce target supervision into coarse structural transfer and oscillatory reconstruction.
- Experiments on SimpleWave, Helmholtz, and Maxwell benchmarks show that APEX consistently improves higher-frequency prediction over direct lower-to-higher extrapolation, target-adapted operator, and joint generative baselines.

## 2 Why Direct Higher-Frequency Extrapolation Fails

### 2.1 Wave Equation and Green's Function

We consider parameterized frequency-domain wave equations of the form

$$\mathcal{L}_{e,\nu}[u](\mathbf{r}) = s(\mathbf{r}; e, \nu), \qquad \mathbf{r} \in \Omega, \tag{1}$$

where $\mathbf{r}$ denotes spatial coordinates, $\Omega$ is the spatial domain, $s(\mathbf{r}; e, \nu)$ is the source term, $e$ denotes the environment and geometry, $\nu$ is the spectral variable, and $u$ is the resulting complex-valued field. Throughout this work, $\nu$ is taken to be positive, corresponding to a physical frequency, wavenumber, or dimensionless spectral parameter depending on the benchmark. This form covers standard frequency-domain wave models, including Maxwell and Helmholtz equations Colton & Kress (2013), as well as the SimpleWave simulator used in this work.

For linear frequency-domain problems, the field admits a Green's-function representation of the form Colton & Kress (2013); Stakgold & Holst (2011)

$$u(\mathbf{r}; e, \nu) = \int_{\Omega} G_{e,\nu}(\mathbf{r}, \mathbf{r}')\, s(\mathbf{r}'; e, \nu)\, d\mathbf{r}', \tag{2}$$

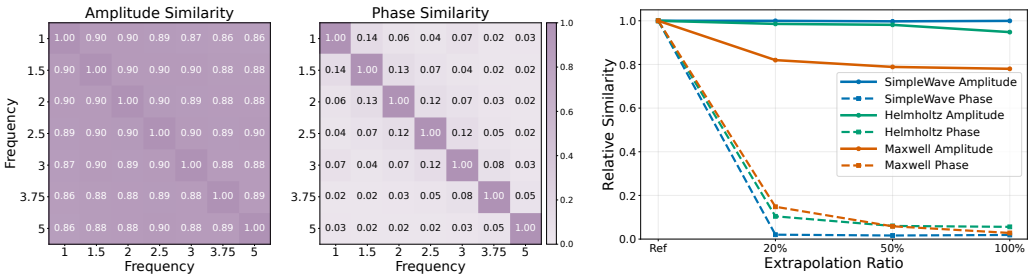


(a) (b) (c)

Figure 1: Empirical evidence for cross-frequency transfer asymmetry. (a,b) Ground-truth amplitude and phase similarity on the Maxwell benchmark. (c) Relative amplitude and phase similarity of lower-frequency FNO extrapolation across the three benchmarks.

where $G_{e,\nu}(\mathbf{r}, \mathbf{r}')$ is the associated Green's function. Eq. (2) is introduced here not to develop a full Green's-function analysis, but to emphasize a structural fact that will matter later: frequency-domain wave fields are organized by source-dependent propagation geometry. This viewpoint motivates why, in the higher-frequency regime, geometric oscillatory cues can be useful even when a full direct surrogate prediction becomes unreliable.

### 2.2 Neural Operators and Direct Lower-to-Higher Extrapolation

Neural operators provide surrogate models for PDE solution operators. A common neural-operator layer can be written as Kovachki et al. (2023); Li et al. (2021)

$$v_{\ell+1}(\mathbf{r}) = \sigma\left(W_\ell v_\ell(\mathbf{r}) + \int_\Omega \kappa_\ell(\mathbf{r}, \mathbf{r}'; e, \nu)\, v_\ell(\mathbf{r}')\, d\mathbf{r}'\right), \tag{3}$$

where $v_\ell$ denotes the latent field at layer $\ell$, $W_\ell$ is a linear transformation, $\sigma(\cdot)$ is a nonlinear activation function, and $\kappa_\ell$ parameterizes an integral kernel.

In the cross-frequency setting, let $\mathcal{V}_{\text{train}} = \{\nu_1, \ldots, \nu_m\}$ denote the observed training frequencies, and let a neural operator $F_\theta$ be learned from samples $\{(e, \nu, u(\mathbf{r}; e, \nu)) : \nu \in \mathcal{V}_{\text{train}}\}$. When the trained model is queried at a new frequency $\tilde{\nu} \notin \mathcal{V}_{\text{train}}$,

$$\hat{u}(\mathbf{r}; e, \tilde{\nu}) = F_\theta(e, \tilde{\nu})(\mathbf{r}), \tag{4}$$

the problem becomes *direct lower-to-higher extrapolation* if $\tilde{\nu}$ lies outside the observed frequency range. This is the naive strategy studied in this paper: train on lower frequencies and directly query the learned operator at higher ones. Existing work has mainly focused on interpolation or in-range prediction Gu et al. (2022); Seo et al. (2024), whereas this out-of-range higher-frequency regime remains much less explored. The key question is therefore not only whether the prediction degrades, but also *how* that degradation is distributed across the complex field representation.

### 2.3 The Asymmetry of Cross-Frequency Transferability

We begin from an empirical observation: cross-frequency variation does not affect all components of a complex wave field equally. To visualize this effect, we first compare the same Maxwell sample across different frequencies using separate similarity measures for amplitude and phase. Let $A(\mathbf{r}; \nu)$ and $\phi(\mathbf{r}; \nu)$ denote the amplitude and phase of the field, respectively. For two frequencies $\nu_i$ and $\nu_j$, we measure amplitude similarity by

$$S_A(\nu_i, \nu_j) = \frac{\langle A(\mathbf{r}; \nu_i), A(\mathbf{r}; \nu_j)\rangle}{\|A(\mathbf{r}; \nu_i)\|_2 \, \|A(\mathbf{r}; \nu_j)\|_2}, \tag{5}$$

and phase similarity by the phase-factor coherence

$$S_P(\nu_i, \nu_j) = \left| \frac{1}{N} \sum_{\mathbf{r}\in\Omega} \exp\big(i\big(\phi(\mathbf{r}; \nu_i) - \phi(\mathbf{r}; \nu_j)\big)\big) \right|. \tag{6}$$

The resulting ground-truth similarity maps are shown in Fig. 1(a)–(b). They reveal a clear asymmetry: amplitude remains comparatively structured across frequency pairs, whereas phase coherence decays much more rapidly as the frequency gap increases.

We next ask whether the same asymmetry also appears in learned prediction. We train an FNO on lower frequencies and evaluate it at unseen higher frequencies. Fig. 1(c) shows relative amplitude and phase similarity across SimpleWave, Helmholtz, and Maxwell benchmarks, obtained by normalizing each higher-frequency prediction similarity by that at the highest observed lower frequency. Benchmark details are provided in Appendix B. In all three cases, amplitude similarity remains substantially higher than phase similarity as the target frequency moves away from the observed regime. Thus, direct lower-to-higher extrapolation does not fail uniformly across the complex field: the phase-sensitive component breaks down earlier and more severely than the amplitude component.

The empirical evidence above suggests that the central difficulty of higher-frequency prediction is not merely larger error, but uneven degradation across different components of the complex field. To formalize this asymmetry, we adopt a local amplitude–phase abstraction. Let $u$ and $\hat{u}$ denote the true and predicted fields, respectively, and write

$$u(\mathbf{r},\nu) = A(\mathbf{r};\nu)\exp\big(i\phi(\mathbf{r};\nu)\big), \qquad \hat{u}(\mathbf{r},\nu) = \hat{A}(\mathbf{r};\nu)\exp\big(i\hat{\phi}(\mathbf{r};\nu)\big). \tag{7}$$

In the local abstraction below, we express phase through effective travel-time functions, i.e., $\phi(\mathbf{r};\nu) \approx \nu\tau(\mathbf{r})$ and $\hat{\phi}(\mathbf{r};\nu) \approx \nu\hat{\tau}(\mathbf{r})$, and define

$$\Delta\tau(\mathbf{r}) := \hat{\tau}(\mathbf{r}) - \tau(\mathbf{r}). \tag{8}$$

Since our interest is in whole-field discrepancy rather than a single spatial point, we measure the error over a measurable region $S \subseteq \Omega$ by

$$\|\hat{u} - u\|^2_{L^2(S)} = \int_S |\hat{u}(\mathbf{r},\nu) - u(\mathbf{r},\nu)|^2 \, d\mathbf{r}. \tag{9}$$

This regional quantity is more natural than a pointwise error here because both the empirical observations in Fig. 1 and the prediction task itself concern the field as a whole. The result below follows by integrating the corresponding pointwise amplitude–phase decomposition over $S$; the pointwise form and proof are given in Appendix A.

**Proposition 1** (Regional decomposition of complex-field error)**.** *For any measurable region $S \subseteq \Omega$,*

$$\|\hat{u} - u\|^2_{L^2(S)} = \|\hat{A} - A\|^2_{L^2(S)} + 4\int_S A(\mathbf{r};\nu)\hat{A}(\mathbf{r};\nu)\,\sin^2\left(\frac{\nu\Delta\tau(\mathbf{r})}{2}\right) d\mathbf{r}. \tag{10}$$

Proposition 1 makes the asymmetry explicit at the level of the regional field error. Amplitude mismatch contributes through a direct term, whereas phase-sensitive mismatch contributes through an oscillatory term whose argument is scaled by frequency. The two components therefore enter higher-frequency prediction error through qualitatively different mechanisms.

The following corollary further highlights the frequency dependence through a simple upper bound.

**Corollary 2** (Frequency-sensitive upper bound for the regional field error)**.** *If $A(\mathbf{r};\nu) \le A_{\max}$ and $\hat{A}(\mathbf{r};\nu) \le \hat{A}_{\max}$ on $S$, then*

$$\|\hat{u} - u\|^2_{L^2(S)} \le \|\hat{A} - A\|^2_{L^2(S)} + \nu^2 A_{\max}\hat{A}_{\max}\,\|\Delta\tau\|^2_{L^2(S)}. \tag{11}$$

Corollary 2 is explanatory rather than tight: even with controlled amplitude mismatch, phase-sensitive propagation mismatch carries an explicit frequency-dependent factor and can grow rapidly in the higher-frequency regime.

**Design Guideline.** Cross-frequency transfer in complex wave fields is inherently asymmetric: coarse amplitude remains comparatively transferable across frequencies, whereas phase-sensitive oscillatory detail is substantially more fragile outside the observed regime. Accordingly, higher-frequency prediction should reuse transferable coarse structure and separately recover the missing oscillatory detail.

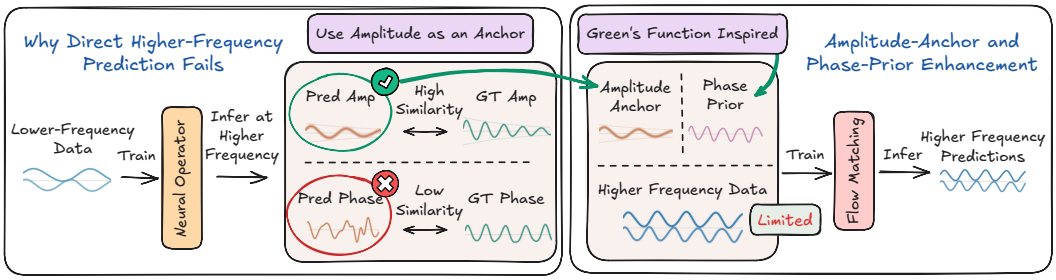


Figure 2: Overview of the proposed framework. Left: direct higher-frequency prediction by a lower-frequency neural operator preserves amplitude structure more reliably than phase. Right: the coarse amplitude is retained as an anchor and combined with a Green's-function-inspired phase prior to guide conditional flow matching.

# 3 METHOD

## 3.1 PROBLEM SETUP

We consider higher-frequency prediction under scarce target supervision for parameterized wave fields. Let $e \in \mathcal{E}$ denote an environment, $\nu \in \mathcal{V}$ denote the spectral variable, and $u(\mathbf{r}; e, \nu) \in \mathbb{C}^{H \times W}$ denote the target complex-valued field on a spatial grid $\mathbf{r}$. We partition the spectrum into a lower-frequency regime $\mathcal{V}_{\mathrm{LF}}$ and a disjoint higher-frequency regime $\mathcal{V}_{\mathrm{HF}}$, and aim to predict $u(\mathbf{r}; e, \nu)$ for $\nu \in \mathcal{V}_{\mathrm{HF}}$ using supervision mainly in $\mathcal{V}_{\mathrm{LF}}$ together with limited target-frequency samples.

Motivated by the transfer asymmetry established in Sec. 2.3, our design separates higher-frequency prediction into two parts: reusing the comparatively transferable coarse structure from a frozen lower-frequency backbone, and recovering the missing oscillatory detail using a conditional enhancement model. The overall pipeline is illustrated in Fig. 2.

## 3.2 COARSE AMPLITUDE ANCHOR

We first extract the transferable component of lower-to-higher prediction. A neural operator $F_\theta$ is trained on the lower-frequency regime $\mathcal{V}_{\mathrm{LF}}$, frozen after training, and then evaluated at target frequencies $\nu \in \mathcal{V}_{\mathrm{HF}}$. Its output is treated as a coarse intermediate prediction rather than as the final higher-frequency solution. We define the corresponding coarse amplitude anchor by

$$a_{\mathrm{coarse}}(\mathbf{r}; e, \nu) := \log\left(\left|F_\theta(e, \nu)(\mathbf{r})\right| + \epsilon\right) \in \mathbb{R}^{H \times W}. \tag{12}$$

Here, $\epsilon$ is a small positive constant for numerical stability. We retain only the log-amplitude of the extrapolated prediction as a coarse target-regime structural reference, leaving phase-sensitive oscillations to subsequent refinement. In this work, $F_\theta$ is instantiated with FNO for simplicity, but the amplitude-anchor construction is backbone-agnostic in principle and can be built on any surrogate that provides a coarse target-frequency prediction.

## 3.3 GREEN'S-FUNCTION-INSPIRED PHASE PRIOR

The coarse amplitude anchor captures transferable large-scale structure, but it does not by itself provide the oscillatory phase information required in the higher-frequency regime. We therefore introduce an explicit phase prior that supplies a lightweight physical guide for the missing phase-sensitive detail. Its role is not to reproduce the full target field, but to provide a compact oscillatory scaffold consistent with the source, the geometry, and an effective phase scale.

Our construction is based on a simple geometric-propagation view of phase: when a small number of dominant propagation paths govern the main oscillatory behavior, the accumulated phase along each path can be approximated by an effective phase coefficient multiplied by the corresponding propagation distance. Based on this intuition, we use the analytic surrogate

$$G_{\mathrm{prior}}(\mathbf{r}, \mathbf{r}_s; e, \nu) = \sum_{m=1}^{M} a_m \exp\big(i\, \kappa_{\mathrm{ref}}(e, \nu)\, L_m(\mathbf{r}, \mathbf{r}_s)\big), \tag{13}$$

where $\mathbf{r}_s$ denotes the source location, $M$ is the number of retained dominant propagation paths, $L_m(\mathbf{r}, \mathbf{r}_s)$ is the geometric path length from the source to location $\mathbf{r}$, $\kappa_{\mathrm{ref}}(e, \nu)$ is a domain-adaptive reference phase coefficient, and $a_m$ is a fixed coefficient associated with the $m$-th retained path. Eq. (13) is a unified low-complexity surrogate for phase accumulation, rather than a full approximation of the Green's function or the target field. The retained paths are therefore chosen to capture the dominant propagation geometry rather than to enumerate all physical paths. Concrete instantiations for the experimental domains are provided in Appendix D.

Since transferable amplitude is already supplied by Sec. 3.2, we use only the surrogate phase:

$$\phi_{\mathrm{base}}(\mathbf{r}; e, \nu) := \mathrm{Arg}\big(G_{\mathrm{prior}}(\mathbf{r}, \mathbf{r}_s; e, \nu)\big). \tag{14}$$

Here, $\mathrm{Arg}(\cdot)$ denotes the principal argument, with values in $(-\pi, \pi]$. This phase map provides source-aware and geometry-aware oscillatory cues, while the remaining higher-frequency detail is recovered by the enhancement model.

### 3.4 Amplitude- and Phase-Prior-Guided Conditional Flow Matching

We use the coarse amplitude anchor and the phase prior as complementary conditions for a generative enhancement model: the former provides transferable large-scale structure, while the latter supplies source- and geometry-aligned oscillatory cues.

Let $\phi = \mathrm{Arg}(u)$ denote the target phase. We represent the target field by

$$x_1 = \big[\log(|u| + \epsilon),\ \sin\phi,\ \cos\phi\big], \tag{15}$$

where the logarithm compresses the dynamic range of the amplitude and the sine–cosine representation avoids the discontinuity of wrapped phase at $\pm\pi$. The conditioning variables are

$$c = \big(a_{\mathrm{coarse}},\ [\sin\phi_{\mathrm{base}}, \cos\phi_{\mathrm{base}}],\ z_f\big), \tag{16}$$

where $a_{\mathrm{coarse}}$ is the coarse amplitude anchor from Sec. 3.2, $[\sin\phi_{\mathrm{base}}, \cos\phi_{\mathrm{base}}]$ are the phase-prior features from Sec. 3.3, and $z_f$ is a Fourier feature embedding of the spectral variable $\nu$.

The conditional velocity field is parameterized by a U-Net Ronneberger et al. (2015). The coarse amplitude and phase-prior maps are injected through multi-scale spatial conditioning so that both global field organization and local oscillatory cues remain available across resolution levels, while $z_f$ modulates the network through FiLM-style conditioning Perez et al. (2018) to adapt the feature response to different target frequencies.

Conditional flow matching (CFM) is used to model the enhancement process Lipman et al. (2023); Tong et al. (2024). With $x_0 \sim \mathcal{N}(0, I)$ and the linear interpolation

$$x_t = (1 - t)x_0 + tx_1, \tag{17}$$

the network $v_\psi(x_t, t; c)$ is trained with

$$\mathcal{L}_{\mathrm{FM}} = \mathbb{E}_{x_0, x_1, t}\left[\|v_\psi(x_t, t; c) - (x_1 - x_0)\|_2^2\right]. \tag{18}$$

At inference, we follow the standard CFM sampling procedure conditioned on $c$.

## 4 Related Work

**Neural operators** have become a standard framework for learning parametric PDE solution maps in function space, with representative formulations including DeepONet and FNO Lu et al. (2021); Li et al. (2021); Kovachki et al. (2023). Subsequent variants improve these models through transformer-based, geometry-aware, and multiscale designs Cao (2021); Hao et al. (2023); Li et al. (2023); Rahman et al. (2023); Cao et al. (2024). Recent work has also begun to study cross-frequency prediction in wave problems. Gu et al. (2022); Seo et al. (2024) consider cross-wavelength prediction within an observed spectral regime. Sun et al. (2025); Wang et al. (2024) show that cross-frequency generalization can be improved through adaptation or fine-tuning on the target regime. (Benitez et al., 2024; Song & Wang, 2022) further examine out-of-distribution behavior and supervised higher-frequency wavefield extrapolation. Compared with this line of work, our study focuses on what fails to transfer across frequencies in oscillatory wave fields and how that failure should be handled.

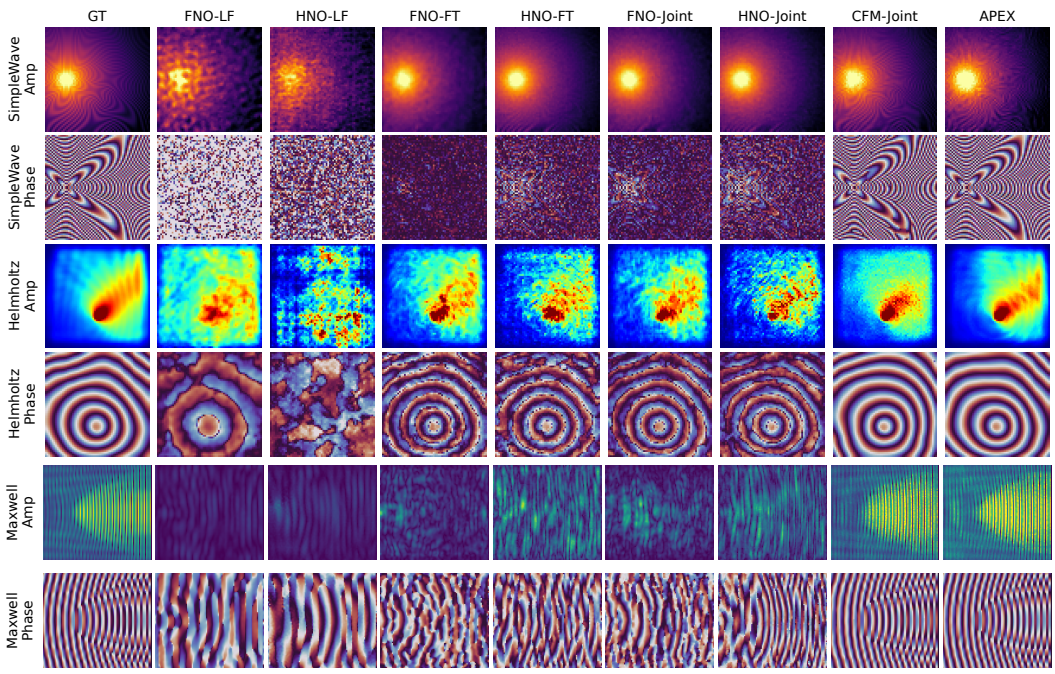


Figure 3: Qualitative comparison at HF100, the largest frequency gap considered in this paper. Across three benchmarks, APEX recovers higher-frequency amplitude and phase patterns with fewer structural distortions and better alignment with the ground truth.

**Generative surrogate models** (Ho et al., 2020; Song et al., 2021) have become a useful tool for scientific field modeling. Beyond standard diffusion formulations, dynamics-aware generative models have been used for long-horizon PDE refinement, and spatiotemporal physical prediction (Lippe et al., 2023; Rühling Cachay et al., 2023). Recent studies further extend generative modeling to physics-aware field reconstruction and PDE-governed scientific simulation (Bastek et al., 2025; Huang et al., 2024; Li et al., 2025; Chen et al., 2025; Zou et al., 2026). In contrast, our work targets scarce-supervision higher-frequency wave prediction, where the generative model is used as an enhancer guided by transferable amplitude structure and a phase-sensitive oscillatory prior.

## 5 Experiments

**Experimental setup.** The evaluation is conducted on three wave benchmarks: SimpleWave, Helmholtz, and Maxwell. SimpleWave and Helmholtz use scalar complex fields on fixed $64 \times 64$ grids, while Maxwell uses the dominant polarization component $E_y$ on an $80 \times 92$ grid. All benchmarks use the same frequency-split protocol: four lower-frequency (LF) values and three **higher-frequency (HF)** values, with $N$ samples per frequency. The $4N$ LF samples are split $80\%/20\%$ for training/testing, and the $3N$ HF samples are split $20\%/80\%$. Evaluation is performed only on the HF test split; the LF test split is used only for lower-frequency monitoring and is not included in the reported HF metrics. **HF20**, **HF50**, and **HF100** denote target groups whose frequencies exceed the maximum LF training frequency by $20\%$, $50\%$, and $100\%$, respectively. Additional details are in Appendix B.

The comparison includes three baseline model families: the Fourier neural operator (FNO) (Li et al., 2021), the Helmholtz neural operator (HNO) (Zou et al., 2024), and conditional flow matching (CFM) (Lipman et al., 2023; Tong et al., 2024). For the operator baselines, we consider three training protocols. 'LF' denotes LF-only training followed by direct extrapolation to the target frequencies. 'FT' initializes from the LF model and then fine-tunes on the available HF training data. 'Joint' denotes joint training on all available LF and HF training samples. This yields FNO-LF/FT/Joint and HNO-LF/FT/Joint as operator baselines, while CFM-Joint is the generative baseline trained jointly on available LF and HF data. APEX first trains and freezes the LF operator on LF samples, then trains the enhancer using the same HF training split.

Table 1: Higher-frequency prediction results across three benchmarks. Each frequency group reports $H^1$/AWPC. Underlined entries denote the best result within each dataset and metric column.

**SimpleWave**

| Method | HF20 | | HF50 | | HF100 | | Overall | |
|---|---|---|---|---|---|---|---|---|
| | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ |
| FNO-LF | 1.377 | 0.012 | 1.454 | 0.013 | 1.429 | 0.013 | 1.420 | 0.013 |
| FNO-FT | 1.053 | 0.157 | 1.114 | 0.238 | 1.128 | 0.133 | 1.099 | 0.176 |
| HNO-LF | 1.359 | 0.012 | 1.420 | 0.014 | 1.401 | 0.014 | 1.393 | 0.013 |
| HNO-FT | 0.923 | 0.496 | 1.021 | 0.465 | 1.071 | 0.400 | 1.005 | 0.454 |
| FNO-Joint | 0.963 | 0.403 | 1.032 | 0.438 | 1.100 | 0.324 | 1.032 | 0.388 |
| HNO-Joint | 0.885 | 0.540 | 0.973 | 0.531 | 1.044 | 0.433 | 0.967 | 0.501 |
| CFM-Joint | 0.935 | 0.620 | 1.151 | 0.516 | 1.234 | 0.351 | 1.107 | 0.496 |
| APEX | 0.396 | 0.909 | 0.494 | 0.865 | 0.804 | 0.659 | 0.564 | 0.811 |

**Helmholtz**

| Method | HF20 | | HF50 | | HF100 | | Overall | |
|---|---|---|---|---|---|---|---|---|
| | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ |
| FNO-LF | 1.611 | 0.169 | 1.686 | 0.098 | 2.261 | 0.050 | 1.853 | 0.106 |
| FNO-FT | 0.762 | 0.897 | 0.783 | 0.874 | 0.966 | 0.759 | 0.837 | 0.843 |
| HNO-LF | 2.989 | 0.097 | 3.743 | 0.042 | 5.637 | 0.032 | 4.123 | 0.057 |
| HNO-FT | 0.756 | 0.871 | 0.909 | 0.821 | 1.276 | 0.696 | 0.980 | 0.796 |
| FNO-Joint | 0.715 | 0.903 | 0.779 | 0.850 | 0.899 | 0.737 | 0.797 | 0.830 |
| HNO-Joint | 0.930 | 0.853 | 1.006 | 0.794 | 1.401 | 0.662 | 1.112 | 0.770 |
| CFM-Joint | 0.852 | 0.653 | 0.895 | 0.543 | 0.930 | 0.397 | 0.893 | 0.531 |
| APEX | 0.422 | 0.928 | 0.492 | 0.885 | 0.634 | 0.815 | 0.516 | 0.876 |

**Maxwell**

| Method | HF20 | | HF50 | | HF100 | | Overall | |
|---|---|---|---|---|---|---|---|---|
| | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ |
| FNO-LF | 1.090 | 0.066 | 1.023 | 0.003 | 1.008 | 0.000 | 1.024 | 0.017 |
| FNO-FT | 0.840 | 0.465 | 0.899 | 0.301 | 0.954 | 0.169 | 0.925 | 0.284 |
| HNO-LF | 1.097 | 0.044 | 1.045 | 0.012 | 1.014 | 0.002 | 1.034 | 0.007 |
| HNO-FT | 0.912 | 0.421 | 0.947 | 0.329 | 0.962 | 0.277 | 0.951 | 0.329 |
| FNO-Joint | 0.904 | 0.422 | 0.933 | 0.337 | 0.959 | 0.233 | 0.945 | 0.313 |
| HNO-Joint | 0.842 | 0.441 | 0.900 | 0.342 | 0.948 | 0.285 | 0.922 | 0.342 |
| CFM-Joint | 0.701 | 0.828 | 0.897 | 0.680 | 1.187 | 0.411 | 1.061 | 0.595 |
| APEX | 0.697 | 0.812 | 0.790 | 0.725 | 0.929 | 0.624 | 0.865 | 0.703 |

Performance is evaluated with two metrics. We use relative complex $H^1$ error (Qiu et al., 2024)

$$H^1(\hat{u}, u) = \left( \frac{\|\hat{u} - u\|_2^2 + \|\Delta_x(\hat{u} - u)\|_2^2 + \|\Delta_y(\hat{u} - u)\|_2^2}{\|u\|_2^2 + \|\Delta_x u\|_2^2 + \|\Delta_y u\|_2^2} \right)^{1/2}, \tag{19}$$

where $\hat{u}$ and $u$ denote the predicted and ground-truth complex wave fields, and $\Delta_x, \Delta_y$ are discrete forward differences along the spatial directions. Amplitude-weighted phase coherence (AWPC) is

$$\mathrm{AWPC}(\hat{u}, u) = \left| \frac{\sum_{\mathbf{x}} |u(\mathbf{x})| \exp\big(i(\mathrm{Arg}(\hat{u}(\mathbf{x})) - \mathrm{Arg}(u(\mathbf{x})))\big)}{\sum_{\mathbf{x}} |u(\mathbf{x})|} \right|, \tag{20}$$

which evaluates phase agreement while weighting each spatial location by the magnitude of the ground-truth field Barnes et al. (2024); Cardenas et al. (2026). All reported results are averaged over the corresponding test set. Additional uncertainty estimates for the main results are in Appendix F.

**Main results.** The proposed method achieves the strongest performance across the three benchmarks. As shown in Fig. 3, direct LF extrapolation retains some coarse amplitude structure but

Table 2: Ablation results on the Helmholtz benchmark.

| Method | HF20 | | HF50 | | HF100 | | Overall | |
|---|---|---|---|---|---|---|---|---|
| | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ |
| w/o Phase Prior | 0.483 | 0.880 | 0.580 | 0.809 | 0.812 | 0.689 | 0.625 | 0.793 |
| w/o Coarse Anchor | 0.907 | 0.628 | 1.081 | 0.505 | 1.252 | 0.379 | 1.080 | 0.504 |
| APEX | 0.422 | 0.928 | 0.492 | 0.885 | 0.634 | 0.815 | 0.516 | 0.876 |

Table 3: Sensitivity to target-frequency supervision on the Helmholtz benchmark. Underlines mark the better result between CFM-Joint and APEX at each ratio.

| Ratio | Method | HF20 | | HF50 | | HF100 | | Overall | |
|---|---|---|---|---|---|---|---|---|---|
| | | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ |
| 1/9 | CFM-Joint | 0.859 | 0.655 | 0.912 | 0.489 | 0.936 | 0.409 | 0.902 | 0.518 |
| 1/9 | APEX | 0.656 | 0.819 | 0.721 | 0.760 | 0.906 | 0.635 | 0.761 | 0.738 |
| 2/8 | CFM-Joint | 0.852 | 0.653 | 0.895 | 0.543 | 0.930 | 0.397 | 0.893 | 0.531 |
| 2/8 | APEX | 0.422 | 0.928 | 0.492 | 0.885 | 0.634 | 0.815 | 0.516 | 0.876 |
| 3/7 | CFM-Joint | 0.878 | 0.603 | 0.908 | 0.521 | 0.938 | 0.401 | 0.908 | 0.508 |
| 3/7 | APEX | 0.494 | 0.929 | 0.534 | 0.902 | 0.663 | 0.829 | 0.564 | 0.886 |
| 4/6 | CFM-Joint | 0.865 | 0.661 | 0.905 | 0.540 | 0.933 | 0.418 | 0.901 | 0.539 |
| 4/6 | APEX | 0.372 | 0.960 | 0.441 | 0.933 | 0.626 | 0.893 | 0.480 | 0.929 |

exhibits clear phase and oscillatory mismatch. Fine-tuned and jointly trained operator baselines improve the coarse pattern but remain limited under scarce target-frequency supervision, while 'CFM-Joint' can produce plausible structures that are not always aligned with the target field. Additional qualitative comparisons at HF20 and HF50 are provided in Appendix E. Tab. 1 further confirms this trend: APEX achieves the best overall results on all three benchmarks, which supports the combination of transferable coarse amplitude and physics-guided phase priors.

**Ablation study.** The ablation results in Tab. 2 examine the two main design components. Removing the phase prior consistently degrades both $H^1$ and AWPC, suggesting that the coarse amplitude anchor alone does not fully recover the required oscillatory structure. Removing the coarse anchor leads to a larger drop, consistent with its role as the main transferable structural reference. The two components therefore appear complementary: the coarse prediction supports global organization, while the phase prior improves fine oscillatory recovery.

**Sensitivity to target-frequency supervision.** Tab. 3 fixes the lower-frequency split and varies only the training/testing ratio in the higher-frequency regime, isolating the effect of target-frequency supervision. Across different higher-frequency supervision ratios, APEX remains consistently better than 'CFM-Joint', indicating that its advantage is not tied to a particular amount of target-frequency data. One possible reason is that, in joint generative training, the lower-frequency samples still dominate the training distribution, making it difficult for 'CFM-Joint' to sufficiently adapt to the scarce higher-frequency regime. This is consistent with Fig. 3: 'CFM-Joint' can generate visually plausible structures, but these structures are not always aligned with the target field. Therefore, the main limitation is not only the amount of target-frequency data, but also the lack of explicit constraints that align the generated field with the target higher-frequency response. APEX addresses this by using transferable structure and phase prior to guide the generated field toward the target response.

## 6 Conclusion

In this work, we identify an asymmetric transfer pattern: coarse amplitude structure remains relatively stable, whereas phase-sensitive oscillatory structure deteriorates rapidly under lower-to-higher extrapolation. APEX exploits this asymmetry with an amplitude anchor and a phase prior, improving prediction over the baselines across three wave benchmarks. A limitation is that the phase prior is simplified; future work will incorporate stronger physics-aware priors and constraints for more accurate and physically consistent higher-frequency prediction.

## AI USE STATEMENT

Generative AI tools were used to assist with writing and language polishing. All AI-assisted edits were reviewed and verified by the authors, who take full responsibility for the final content of this paper.

## REFERENCES


Samuel JK Barnes, J Bjerkan, PT Clemson, J Newman, and A Stefanovska. Phase coherence—a time-localized approach to studying interactions. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 34(7), 2024.

Jan-Hendrik Bastek, WaiChing Sun, and Dennis Kochmann. Physics-informed diffusion models. In *International Conference on Learning Representations*, volume 2025, pp. 3360–3385, 2025.

Jose Antonio Lara Benitez, Takashi Furuya, Florian Faucher, Anastasis Kratsios, Xavier Tricoche, and Maarten V de Hoop. Out-of-distributional risk bounds for neural operators with applications to the helmholtz equation. *Journal of Computational Physics*, 513:113168, 2024.

Qianying Cao, Somdatta Goswami, and George Em Karniadakis. Laplace neural operator for solving differential equations. *Nature Machine Intelligence*, 6(6):631–640, 2024.

Shuhao Cao. Choose a transformer: Fourier or galerkin. *Advances in neural information processing systems*, 34:24924–24940, 2021.

Araceli R Cardenas, Juan F Ramirez-Villegas, Christopher K Kovach, Phillip E Gander, Rachel C Cole, Andrew J Grossbach, Hiroto Kawasaki, Jeremy DW Greenlee, Matthew A Howard, Kirill V Nourski, et al. Exercise enhances hippocampal-cortical ripple interactions in the human brain. *Brain Communications*, 8(2):fcag041, 2026.

Panqi Chen, Yifan Sun, Lei Cheng, Yang Yang, Weichang Li, Yang Liu, Weiqing Liu, Jiang Bian, and Shikai Fang. Generating full-field evolution of physical dynamics from irregular sparse observations. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

David Colton and Rainer Kress. *Inverse Acoustic and Electromagnetic Scattering Theory*. Springer, 3 edition, 2013.

Jiaqi Gu, Zhengqi Gao, Chenghao Feng, Hanqing Zhu, Ray Chen, Duane Boning, and David Pan. Neurolight: A physics-agnostic neural operator enabling parametric photonic device simulation. *Advances in Neural Information Processing Systems*, 35:14623–14636, 2022.

Zhongkai Hao, Zhengyi Wang, Hang Su, Chengyang Ying, Yinpeng Dong, Songming Liu, Ze Cheng, Jian Song, and Jun Zhu. Gnot: A general neural operator transformer for operator learning. In *International conference on machine learning*, pp. 12556–12569. PMLR, 2023.

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.

Jiahe Huang, Guandao Yang, Zichen Wang, and Jeong Joon Park. Diffusionpde: Generative pde-solving under partial observation. *Advances in Neural Information Processing Systems*, 37: 130291–130323, 2024.

F. Ihlenburg and I. Babuška. Finite element solution of the helmholtz equation with high wave number part i: The h-version of the fem. *Computers & Mathematics with Applications*, 30(9):9–37, 1995. ISSN 0898-1221.

Joseph B. Keller. Geometrical theory of diffraction. *J. Opt. Soc. Am.*, 52(2):116–130, Feb 1962.

Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Neural operator: Learning maps between function spaces with applications to pdes. *Journal of Machine Learning Research*, 24(89):1–97, 2023.

Zeyu Li, Hongkun Dou, Shen Fang, Wang Han, Yue Deng, and Lijun Yang. Physics-aligned field reconstruction with diffusion bridge. In *The Thirteenth International Conference on Learning Representations*, 2025.

Zongyi Li, Nikola Borislavov Kovachki, Kamyar Azizzadenesheli, Kaushik Bhattacharya, Andrew Stuart, Anima Anandkumar, et al. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021.

Zongyi Li, Nikola Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Otta, Mohammad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, et al. Geometry-informed neural operator for large-scale 3d pdes. *Advances in Neural Information Processing Systems*, 36: 35836–35854, 2023.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations*, 2023.

Phillip Lippe, Bas Veeling, Paris Perdikaris, Richard Turner, and Johannes Brandstetter. Pde-refiner: Achieving accurate long rollouts with neural pde solvers. *Advances in Neural Information Processing Systems*, 36:67398–67433, 2023.

Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deeponet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3):218–229, 2021.

Ethan Perez, Florian Strub, Harm De Vries, Vincent Dumoulin, and Aaron Courville. Film: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.

Yuan Qiu, Nolan Bridges, and Peng Chen. Derivative-enhanced deep operator network. *Advances in Neural Information Processing Systems*, 37:20945–20981, 2024.

Md Ashiqur Rahman, Zachary E Ross, and Kamyar Azizzadenesheli. U-NO: U-shaped neural operators. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856.

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241. Springer, 2015.

Salva Rühling Cachay, Bo Zhao, Hailey Joren, and Rose Yu. Dyffusion: A dynamics-informed diffusion model for spatiotemporal forecasting. *Advances in neural information processing systems*, 36:45259–45287, 2023.

Joonhyuk Seo, Chanik Kang, Dongjin Seo, and Haejun Chung. Wave interpolation neural operator: Interpolated prediction of electric fields across untrained wavelengths. In *NeurIPS 2024 Workshop on Data-driven and Differentiable Simulations, Surrogates, and Solvers*, 2024.

J A Sethian. A fast marching level set method for monotonically advancing fronts. *Proceedings of the National Academy of Sciences*, 93(4):1591–1595, 1996.

Chao Song and Yanghua Wang. High-frequency wavefield extrapolation using the fourier neural operator. *Journal of Geophysics and Engineering*, 19(2):269–282, 2022.

Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *International Conference on Learning Representations*, 2021.

Ivar Stakgold and Michael J. Holst. *Green's functions and boundary value problems*. Wiley, 3 edition, 2011.

Yifan Sun, Lei Cheng, Jianlong Li, and Peter Gerstoft. Hankel-fno: Fast underwater acoustic charting via physics-encoded fourier neural operator. *The Journal of the Acoustical Society of America*, 158 (6):5075–5089, 2025.

Alexander Tong, Kilian FATRAS, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856.

Yufeng Wang, Heng Zhang, Chensen Lai, and Xiangyun Hu. Transfer learning fourier neural operator for solving parametric frequency-domain wave equations. *IEEE Transactions on Geoscience and Remote Sensing*, 62:1–11, 2024.

Caifeng Zou, Kamyar Azizzadenesheli, Zachary E Ross, and Robert W Clayton. Deep neural helmholtz operators for 3-d elastic wave propagation and inversion. *Geophysical Journal International*, 239(3):1469–1484, 2024.

Yicheng Zou, Samuel Lanthaler, and Hossein Salahshoor. A probabilistic framework for solving high-frequency helmholtz equations via diffusion models. *arXiv preprint arXiv:2602.04082*, 2026.

# A Proofs for the Amplitude–Phase Error Decomposition

This appendix provides the proofs of the analytical results stated in Sec. 2.3. We first derive an exact pointwise amplitude–phase error decomposition, then integrate it over a spatial region to obtain Proposition 1, and finally establish Corollary 2.

## A.1 Exact Pointwise Decomposition

**Lemma 3** (Exact pointwise amplitude–phase error decomposition)**.** *Fix a target frequency $\nu > 0$. Let*

$$u(\mathbf{r},\nu) = A(\mathbf{r};\nu)\exp\big(i\nu\tau(\mathbf{r})\big), \qquad \hat{u}(\mathbf{r},\nu) = \hat{A}(\mathbf{r};\nu)\exp\big(i\nu\hat{\tau}(\mathbf{r})\big), \tag{21}$$

*and define*

$$\Delta\tau(\mathbf{r}) := \hat{\tau}(\mathbf{r}) - \tau(\mathbf{r}). \tag{22}$$

*Then, for every $\mathbf{r} \in \Omega$,*

$$|\hat{u}(\mathbf{r},\nu) - u(\mathbf{r},\nu)|^2 = \big(\hat{A}(\mathbf{r};\nu) - A(\mathbf{r};\nu)\big)^2 + 4A(\mathbf{r};\nu)\hat{A}(\mathbf{r};\nu)\sin^2\!\left(\frac{\nu\Delta\tau(\mathbf{r})}{2}\right). \tag{23}$$

*Proof.* Fix $\mathbf{r} \in \Omega$ and suppress the dependence on $(\mathbf{r},\nu)$ when no ambiguity arises. By direct expansion,

$$|\hat{u} - u|^2 = \left|\hat{A}e^{i\nu\hat{\tau}} - Ae^{i\nu\tau}\right|^2 \tag{24}$$

$$= \hat{A}^2 + A^2 - 2A\hat{A}\cos\big(\nu(\hat{\tau} - \tau)\big). \tag{25}$$

Using $\Delta\tau = \hat{\tau} - \tau$ and the identity

$$\cos x = 1 - 2\sin^2(x/2), \tag{26}$$

we obtain

$$|\hat{u} - u|^2 = \hat{A}^2 + A^2 - 2A\hat{A}\Big(1 - 2\sin^2(\nu\Delta\tau/2)\Big) \tag{27}$$

$$= (\hat{A} - A)^2 + 4A\hat{A}\sin^2\!\left(\frac{\nu\Delta\tau}{2}\right), \tag{28}$$

which proves (23). □

## A.2 Proof of Proposition 1

*Proof of Proposition 1.* By definition of the regional $L^2$ field error,

$$\|\hat{u} - u\|^2_{L^2(S)} = \int_S |\hat{u}(\mathbf{r},\nu) - u(\mathbf{r},\nu)|^2\, d\mathbf{r}. \tag{29}$$

Applying Lemma 3 pointwise inside the integral yields

$$\|\hat{u} - u\|^2_{L^2(S)} = \int_S \left[\big(\hat{A}(\mathbf{r};\nu) - A(\mathbf{r};\nu)\big)^2 + 4A(\mathbf{r};\nu)\hat{A}(\mathbf{r};\nu)\sin^2\!\left(\frac{\nu\Delta\tau(\mathbf{r})}{2}\right)\right] d\mathbf{r} \tag{30}$$

$$= \|\hat{A} - A\|^2_{L^2(S)} + 4\int_S A(\mathbf{r};\nu)\hat{A}(\mathbf{r};\nu)\sin^2\!\left(\frac{\nu\Delta\tau(\mathbf{r})}{2}\right) d\mathbf{r}, \tag{31}$$

which is exactly (10). □

## A.3 Proof of Corollary 2

*Proof of Corollary 2.* Starting from Proposition 1,

$$\|\hat{u} - u\|^2_{L^2(S)} = \|\hat{A} - A\|^2_{L^2(S)} + 4\int_S A(\mathbf{r};\nu)\hat{A}(\mathbf{r};\nu)\sin^2\!\left(\frac{\nu\Delta\tau(\mathbf{r})}{2}\right) d\mathbf{r}. \tag{32}$$

Using the elementary inequality

$$\sin z \le z, \qquad z \ge 0, \tag{33}$$

we obtain

$$\sin^2\left(\frac{\nu\Delta\tau(\mathbf{r})}{2}\right) = \sin^2\left(\frac{\nu|\Delta\tau(\mathbf{r})|}{2}\right) \le \left(\frac{\nu|\Delta\tau(\mathbf{r})|}{2}\right)^2 = \frac{\nu^2}{4}\Delta\tau(\mathbf{r})^2. \tag{34}$$

Substituting this bound into Proposition 1 gives

$$\|\hat{u} - u\|^2_{L^2(S)} \le \|\hat{A} - A\|^2_{L^2(S)} + 4\int_S A(\mathbf{r};\nu)\hat{A}(\mathbf{r};\nu)\cdot\frac{\nu^2}{4}\Delta\tau(\mathbf{r})^2\,d\mathbf{r} \tag{35}$$

$$= \|\hat{A} - A\|^2_{L^2(S)} + \nu^2\int_S A(\mathbf{r};\nu)\hat{A}(\mathbf{r};\nu)\Delta\tau(\mathbf{r})^2\,d\mathbf{r}, \tag{36}$$

If, in addition, $A(\mathbf{r};\nu) \le A_{\max}$ and $\hat{A}(\mathbf{r};\nu) \le \hat{A}_{\max}$ on $S$, then

$$A(\mathbf{r};\nu)\hat{A}(\mathbf{r};\nu) \le A_{\max}\hat{A}_{\max}, \qquad \forall\,\mathbf{r} \in S. \tag{37}$$

Hence

$$\|\hat{u} - u\|^2_{L^2(S)} \le \|\hat{A} - A\|^2_{L^2(S)} + \nu^2 A_{\max}\hat{A}_{\max}\int_S \Delta\tau(\mathbf{r})^2\,d\mathbf{r} \tag{38}$$

$$= \|\hat{A} - A\|^2_{L^2(S)} + \nu^2 A_{\max}\hat{A}_{\max}\,\|\Delta\tau\|^2_{L^2(S)}, \tag{39}$$

which proves (11). □

### A.4 Remark on Scope

The results above should be interpreted as structural characterizations of the regional field error under the local amplitude–phase abstraction used in Sec. 2.3. Proposition 1 is exact and shows that amplitude mismatch and phase-sensitive mismatch enter the global prediction error through qualitatively different mechanisms. Corollary 2 further makes explicit that the phase-induced component carries a frequency-dependent scaling factor. Their role in the main text is explanatory rather than exhaustive: they formalize one mechanism behind the empirical asymmetry observed in Fig. 1, namely that phase-sensitive mismatch becomes increasingly costly in higher-frequency prediction, while coarse amplitude remains comparatively more stable and reusable across frequencies.

## B Dataset Construction

This appendix provides additional details on the benchmark construction and data splits used in the experiments. The three benchmarks instantiate the same lower-to-higher frequency prediction setting studied in the main text: a lower-frequency regime $\mathcal{V}_{\mathrm{LF}}$ is used to train the operator backbone, while a disjoint higher-frequency regime $\mathcal{V}_{\mathrm{HF}}$ is used for target-frequency adaptation and evaluation. Tables 4 and 5 summarize the corresponding dataset specifications and train/test splits.

**SimpleWave benchmark.** The SimpleWave benchmark is generated by a two-path wave simulator on a two-dimensional square domain $\Omega = [0, 10] \times [0, 10]$ in dimensionless spatial coordinates. The field is discretized on a $64 \times 64$ grid. Each sample is specified by an environment parameter $e = v$, where $v$ is a scalar propagation speed sampled uniformly from $[0.8, 1.2]$, and by a frequency $\nu$, which is treated as a dimensionless spectral parameter in this benchmark. The model input is represented as two spatial channels, $[v, \nu]$, where both quantities are broadcast over the grid. The output is a scalar complex field $u(\mathbf{r}; v, \nu) \in \mathbb{C}^{64\times 64}$.

The simulator combines a direct propagation path and a reflected path. Let $\mathbf{r}_s = (2.0, 5.0)$ denote the source location and let $\mathbf{r}_m = (2.0, -5.0)$ denote the mirror source used to define the reflected path. For a spatial location $\mathbf{r} = (x, y)$, define

$$R_1(\mathbf{r}) = \|\mathbf{r} - \mathbf{r}_s\|_2, \qquad R_2(\mathbf{r}) = \|\mathbf{r} - \mathbf{r}_m\|_2. \tag{40}$$

The corresponding travel-time functions are

$$\tau_1(\mathbf{r}; v) = \frac{R_1(\mathbf{r})}{v}\left(1 + \eta\, q_1(\mathbf{r})\right), \qquad \tau_2(\mathbf{r}; v) = \frac{R_2(\mathbf{r}) + d}{v}\left(1 + \eta\, q_2(\mathbf{r})\right), \tag{41}$$

where $q_1$ and $q_2$ are fixed smooth spatial perturbation fields, $\eta = 0.08$ controls their strength, and $d = 0.35$ is a delay bias for the reflected path. Specifically, with $L_x = L_y = 10$, we use

$$q_1(\mathbf{r}) = 0.45 \sin\left(\frac{2\pi x}{L_x}\right) + 0.35 \cos\left(\frac{2\pi y}{L_y}\right) + 0.20 \sin\left(\frac{2\pi(x + 0.6y)}{L_x}\right), \tag{42}$$

$$q_2(\mathbf{r}) = 0.40 \cos\left(\frac{2\pi(x - 0.3y)}{L_x}\right) + 0.30 \sin\left(\frac{2\pi y}{L_y}\right) + 0.15 \cos\left(\frac{2\pi(x + y)}{L_y}\right). \tag{43}$$

These perturbation fields are fixed across samples. The complex field is then generated as

$$u(\mathbf{r}; v, \nu) = a_1(\mathbf{r}) \exp\left(i\, 2\pi\nu\, \tau_1(\mathbf{r}; v)\right) + \lambda\, a_2(\mathbf{r}) \exp\left(i\, 2\pi\nu\, \tau_2(\mathbf{r}; v)\right), \tag{44}$$

with reflection coefficient $\lambda = 0.18$. Here the factor $2\pi$ reflects the cyclic-frequency convention used in the simulator; in the theoretical notation of Sec. 2.2, this constant factor is absorbed into the generic spectral variable $\nu$. The two spatial amplitude envelopes are fixed deterministic functions of location:

$$a_1(\mathbf{r}) = \frac{\exp(-\alpha_1 R_1(\mathbf{r}))}{(R_1(\mathbf{r}) + c_1)^{p_1}}, \qquad a_2(\mathbf{r}) = \frac{\exp(-\alpha_2 R_2(\mathbf{r}))}{\sqrt{R_2(\mathbf{r}) + \epsilon}}, \tag{45}$$

where $\alpha_1 = 0$, $c_1 = 0.8$, $p_1 = 0.3$, $\alpha_2 = 0.12$, and $\epsilon$ is a small numerical constant. Thus, the amplitude envelopes are not learned and do not vary randomly across samples; sample-to-sample variation is induced by the sampled propagation speed $v$ and by the queried frequency $\nu$, while the path geometry and envelope parameters are fixed.

The lower-frequency regime is

$$\mathcal{V}_{\mathrm{LF}} = \{1, 2, 3, 4\}, \tag{46}$$

and the higher-frequency regime is

$$\mathcal{V}_{\mathrm{HF}} = \{4.8, 6, 8\}, \tag{47}$$

where $\nu$ is treated as a dimensionless spectral parameter. These three higher-frequency values correspond to $20\%$, $50\%$, and $100\%$ extrapolation beyond the largest lower-frequency value, respectively. For each frequency, 250 samples are generated. The lower-frequency data are split into 800 training samples and 200 test samples, while the higher-frequency data are split into 150 training samples and 600 test samples.

**Helmholtz benchmark.** The Helmholtz benchmark is generated by solving a two-dimensional heterogeneous frequency-domain Helmholtz equation on a square domain $\Omega = [0, 1] \times [0, 1]$ in dimensionless spatial coordinates. The field is discretized on a $64 \times 64$ grid. For each sample, the complex field $u(\mathbf{r}; n, k)$ satisfies

$$-\Delta u(\mathbf{r}) - k^2 n(\mathbf{r})\big(1 + i\sigma_{\mathrm{abs}}(\mathbf{r})\big)u(\mathbf{r}) = s(\mathbf{r}). \tag{48}$$

Here $n(\mathbf{r})$ is the heterogeneous medium, $k$ is the wavenumber, $\sigma_{\mathrm{abs}}(\mathbf{r})$ is a boundary sponge profile, and $s(\mathbf{r})$ is a localized complex source. The medium is generated from a smooth Gaussian random field. Let $\xi(\mathbf{r})$ denote a standard Gaussian noise field and let $\mathcal{G}_{\ell_g}$ denote Gaussian smoothing on the grid with smoothing width $\ell_g = 8.0$ grid cells. We form a normalized smooth field

$$g(\mathbf{r}) = \frac{\mathcal{G}_{\ell_g}\xi(\mathbf{r}) - \mu_g}{s_g}, \tag{49}$$

where $\mu_g$ and $s_g$ are the empirical mean and standard deviation after smoothing. The medium is then

$$n(\mathbf{r}) = \mathrm{clip}\left(1 + 0.25\, g(\mathbf{r}),\, 0.6,\, 1.4\right). \tag{50}$$

Thus, the sample-to-sample variability in the Helmholtz benchmark comes from the sampled medium field $n(\mathbf{r})$ and the queried wavenumber $k$.

The source is a localized complex Gaussian centered at $\mathbf{r}_s = (0.30, 0.50)$:

$$s(\mathbf{r}) = A_s \exp\left(-\frac{\|\mathbf{r} - \mathbf{r}_s\|_2^2}{2\sigma_s^2}\right) \exp\left(i\beta\, \frac{(\mathbf{r} - \mathbf{r}_s)^\top \mathbf{d}_s}{\sigma_s}\right), \tag{51}$$

where $A_s = 1$, $\sigma_s = 0.03$, $\beta = \pi/6$, and $\mathbf{d}_s = (\cos(\pi/4), \sin(\pi/4))$. The phase-ramp term introduces a mild directional bias near the source while keeping the source localized.

We apply a sponge profile near the boundary. Let $D_x(\mathbf{r})$ and $D_y(\mathbf{r})$ denote the distances from $\mathbf{r}$ to the nearest vertical and horizontal boundaries, respectively. With absorbing width $w_{\mathrm{abs}} = 0.15$, strength $\gamma_{\mathrm{abs}} = 1.5$, and power $p_{\mathrm{abs}} = 2$, the profile is

$$\sigma_{\mathrm{abs}}(\mathbf{r}) = \gamma_{\mathrm{abs}} \left[ \left( \frac{w_{\mathrm{abs}} - D_x(\mathbf{r})}{w_{\mathrm{abs}}} \right)_+^{p_{\mathrm{abs}}} + \left( \frac{w_{\mathrm{abs}} - D_y(\mathbf{r})}{w_{\mathrm{abs}}} \right)_+^{p_{\mathrm{abs}}} \right], \tag{52}$$

where $(\cdot)_+ = \max(\cdot, 0)$. The equation is solved by a second-order finite-difference discretization, yielding a sparse linear system for the interior grid values.

The model input is represented as two spatial channels, $[n, k]$, where $k$ is broadcast over the grid. The target is a scalar complex field $u(\mathbf{r}; n, k) \in \mathbb{C}^{64\times 64}$. For this benchmark, the generic spectral variable $\nu$ corresponds to the wavenumber $k$ used in the finite-difference solver. The lower-frequency regime is

$$\mathcal{V}_{\mathrm{LF}} = \{10, 15, 20, 25\}, \tag{53}$$

and the higher-frequency regime is

$$\mathcal{V}_{\mathrm{HF}} = \{30, 37.5, 50\}, \tag{54}$$

where $\nu$ is treated as a dimensionless spectral parameter. These three higher-frequency values correspond to $20\%$, $50\%$, and $100\%$ extrapolation beyond the largest lower-frequency value, respectively. For each spectral value, 500 samples are generated. The lower-frequency data are split into 1600 training samples and 400 test samples, while the higher-frequency data are split into 300 training samples and 1200 test samples.

**Maxwell benchmark.** The Maxwell benchmark is generated using full-wave electromagnetic simulations in CST Studio Suite. We employ the time-domain solver and place electric-field monitors at multiple prescribed frequencies to extract the corresponding frequency-domain field responses from a single broadband simulation. For each dielectric configuration, the simulated electric-field distribution is recorded on the observation plane at the selected frequencies, yielding complex-valued field samples for cross-frequency prediction. The physical region spans $50\,\mathrm{cm}$ along each coordinate direction, and the learning problem uses two-dimensional $(y, z)$ observation slices sampled on an $80 \times 92$ grid. Each dielectric configuration consists of five horizontally stacked material layers, with relative permittivity increasing monotonically from the top layer to the bottom layer. The model input is represented by three spatial channels $[\varepsilon_r, x, f]$, where $\varepsilon_r(y, z)$ is the relative-permittivity distribution on the observation slice, $x$ is the slice coordinate broadcast over the grid, and $f$ is the excitation frequency broadcast over the grid.

The field is excited by a localized electric dipole source placed $1.5\,\mathrm{cm}$ above the top material surface and horizontally centered in the domain. Owing to its compact spatial support relative to the computational domain and sampling scale, the dipole is treated as a point-like excitation source. The upper and lateral boundaries use open-space radiation conditions, while the lower boundary uses an electric-wall condition,

$$E_t = 0, \tag{55}$$

where $E_t$ denotes the tangential electric field.

The target is the selected dominant polarization component

$$u(y, z; \varepsilon_r, x, f) = E_y(y, z; f) \in \mathbb{C}^{80\times 92}, \tag{56}$$

extracted from the CST frequency-domain solution. In this benchmark, the generic spectral variable $\nu$ corresponds to the physical frequency $f$ in GHz. The lower-frequency regime is

$$\mathcal{V}_{\mathrm{LF}} = \{1.0, 1.5, 2.0, 2.5\}\,\mathrm{GHz}, \tag{57}$$

and the higher-frequency regime is

$$\mathcal{V}_{\mathrm{HF}} = \{3.0, 3.75, 5.0\}\,\mathrm{GHz}. \tag{58}$$

These higher-frequency values correspond to $20\%$, $50\%$, and $100\%$ extrapolation beyond the largest lower-frequency value of $2.5\,\mathrm{GHz}$, respectively. We use 26 dielectric configurations and 30 observation slices for each frequency, giving 780 samples per frequency. The lower-frequency data are split into 2496 training samples and 624 test samples, while the higher-frequency data are split into 468 training samples and 1872 test samples.

Table 4: Dataset construction and spectral splits used in the benchmark experiments.

| Dataset | Generator | Grid | Input channels | Target | LF spectral set | HF spectral set |
|---|---|---|---|---|---|---|
| SimpleWave | two-path wave simulator | $64 \times 64$ | $[v, \nu]$ | scalar complex field | $\{1, 2, 3, 4\}$ | $\{4.8, 6, 8\}$ |
| Helmholtz | finite-difference Helmholtz solver | $64 \times 64$ | $[n, k]$ | scalar complex field | $\{10, 15, 20, 25\}$ | $\{30, 37.5, 50\}$ |
| Maxwell | time domain solver | $80 \times 92$ | $[\varepsilon_r, x, f]$ | $E_y$ complex field | $\{1.0, 1.5, 2.0, 2.5\}$ GHz | $\{3.0, 3.75, 5.0\}$ GHz |

Table 5: Train/test splits used in the benchmark experiments. The LF split is used for training the frozen operator backbone, while the HF split is used for target-frequency enhancement and evaluation. Joint splits are used by joint-training baselines.

| Dataset | LF total | LF train/test | HF total | HF train/test | Joint train/test |
|---|---|---|---|---|---|
| SimpleWave | $4 \times 250 = 1000$ | $800/200$ | $3 \times 250 = 750$ | $150/600$ | $950/600$ |
| Helmholtz | $4 \times 500 = 2000$ | $1600/400$ | $3 \times 500 = 1500$ | $300/1200$ | $1900/1200$ |
| Maxwell | $4 \times 780 = 3120$ | $2496/624$ | $3 \times 780 = 2340$ | $468/1872$ | $2964/1872$ |

## C Training and Model Hyperparameters

All experiments are conducted on a single NVIDIA RTX 4090 GPU with 24 GB memory. All operator baselines use unit-Gaussian normalization for inputs and outputs, with complex fields represented in magnitude–phase form. The FNO baselines are trained for 200 epochs. SimpleWave uses a 4-layer FNO with 24 Fourier modes, hidden width 48, and lifting/projection width 128. Helmholtz uses a 4-layer FNO with 24 Fourier modes, hidden width 64, and lifting/projection width 256. Maxwell uses a 4-layer FNO with 32 Fourier modes, hidden width 64, and lifting/projection width 256. The HNO baselines are also trained for 200 epochs: SimpleWave uses width 8 with 24 Fourier modes and factor 0.75, Helmholtz uses width 11 with 24 Fourier modes and factor 0.75, and Maxwell uses width 11 with 32 Fourier modes and factor 0.75.

CFM-Joint and the second stage of APEX use the same U-Net architecture: base width 128, channel multipliers $[1, 2, 4]$, two residual blocks per resolution level, mid-block attention with 4 heads, and time embedding dimension 64. The CFM models are trained for 500 epochs on SimpleWave, 700 epochs on Helmholtz, and 1000 epochs on Maxwell. At inference time, all CFM models use midpoint ODE sampling with 50 steps.

CFM-Joint is trained on the joint LF+HF split and conditions on the dataset input channels. For SimpleWave and Helmholtz, these channels contain the environment field and spectral variable. For Maxwell, they contain the permittivity map $\epsilon$, the slice coordinate $x$, and the spectral variable.

APEX uses the same CFM architecture and training epochs, but augments the conditioning with the coarse amplitude anchor $a_{\mathrm{coarse}}$ and sine–cosine phase-prior maps $[\sin \phi_{\mathrm{base}}, \cos \phi_{\mathrm{base}}]$, as defined in the main text. SimpleWave and Helmholtz use these spatial maps together with their environment fields. Maxwell follows the same interface and additionally includes the slice coordinate $x$ in the environment representation.

## D Physics-Aware Instantiations of the Phase Prior

This appendix provides the concrete instantiations of the phase prior in Sec. 3.3 for the experimental domains considered in this paper. The prior is intentionally lightweight: its role is to provide a coarse oscillatory scaffold consistent with source location, geometry, and an effective phase scale, rather than to reproduce the full target solution. In particular, the prior is used only through its angle in the main model, while transferable amplitude information is supplied separately by the coarse anchor.

The concrete instantiations used in the three benchmarks are summarized in Tab. 6. For SimpleWave and Helmholtz, the prior retains only the dominant direct path, even though the SimpleWave data generator contains an additional reflected component. This choice keeps the prior deliberately simple and avoids reproducing the full simulator. For Maxwell, we retain a direct path together with one dominant reflected / mirrored path to capture the main geometric oscillatory pattern on the observation slice.

The quantities $v_{\mathrm{ref}}(e)$, $n_{\mathrm{ref}}(e)$, and $\epsilon_{\mathrm{ref}}(e)$ denote sample-wise effective propagation speed, Helmholtz medium coefficient, and permittivity, respectively. Here $L_m(\mathbf{r}, \mathbf{r}_s)$ denotes the geometric path length of the $m$-th retained path from source $\mathbf{r}_s$ to location $\mathbf{r}$. The carrier coefficient $k_0$ converts the benchmark-specific spectral variable into an effective path-accumulated phase scale, and $\kappa_{\mathrm{ref}}(e, \nu)$ then adapts this scale to the current sample through the corresponding medium summary.

Table 6: Physics-aware instantiations of the phase prior used in the experimental domains.

| Domain | Environment | $G_{\mathrm{prior}}$ | $\kappa_{\mathrm{ref}}(e, \nu)$ | $M$ | $a_m$ |
|---|---|---|---|---|---|
| SimpleWave | speed $v$ | $\exp(i\,\kappa_{\mathrm{ref}} L_1)$ | $k_0\, v_{\mathrm{ref}}^{-1}(e)$ | 1 | $a_1 = 1$ |
| Helmholtz | medium $n$ | $\exp(i\,\kappa_{\mathrm{ref}} L_1)$ | $k_0\, \sqrt{n_{\mathrm{ref}}(e)}$ | 1 | $a_1 = 1$ |
| Maxwell | medium $\epsilon$ | $a_1\, \exp(i\,\kappa_{\mathrm{ref}} L_1) + a_2\, \exp(i\,\kappa_{\mathrm{ref}} L_2)$ | $k_0\, \sqrt{\epsilon_{\mathrm{ref}}(e)}$ | 2 | $a_1 = 1,\ a_2 = 1$ |

## E Additional Qualitative Results

Figs. 4 and 5 complement the HF100 comparison in Fig. 3 with smaller frequency gaps. At HF20, most baselines recover the coarse field layout more clearly than at HF100, but they still exhibit visible phase misalignment and local structural errors. APEX better preserves the dominant amplitude organization while maintaining more coherent phase patterns across the three benchmarks.

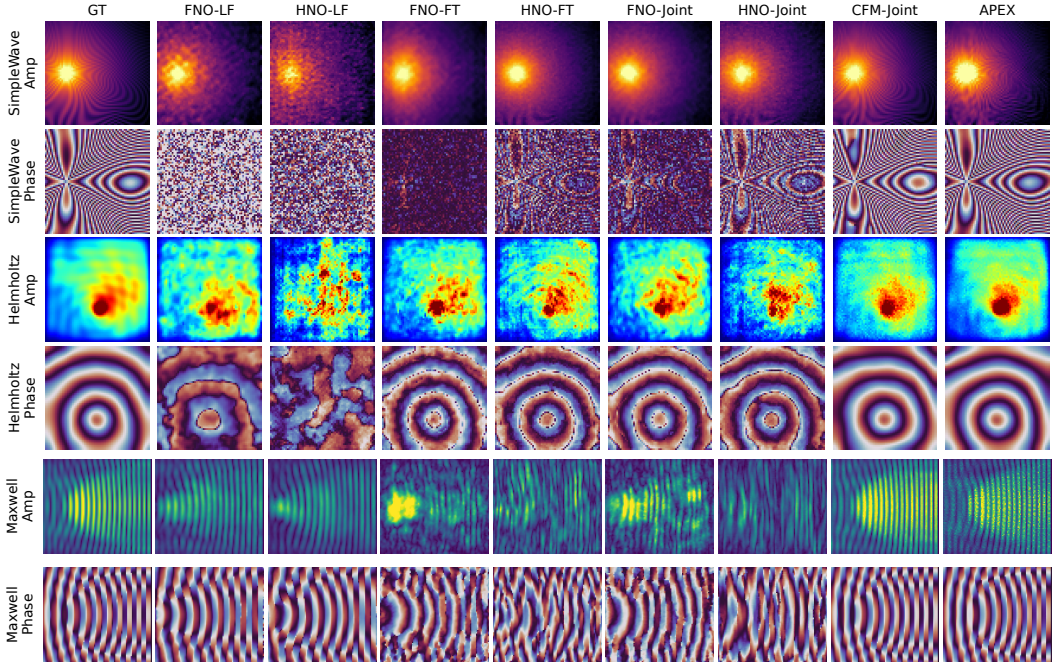


Figure 4: Additional qualitative comparison at HF20.

At HF50, the gap from the lower-frequency regime is larger and the difference between methods becomes more apparent. Direct extrapolation and jointly trained baselines often preserve plausible low-frequency structure but lose fine oscillatory alignment, whereas APEX remains closer to the ground truth in both amplitude localization and phase organization.

## F Bootstrap Confidence Intervals for Main Results

Table 7 complements Tab. 1 by reporting uncertainty estimates for the higher-frequency prediction results. Each entry is shown as mean $\pm$ half-width of a 95% non-parametric bootstrap confidence interval over the fixed HF test split, using $B$=2000 resamples. Following the main text, we report

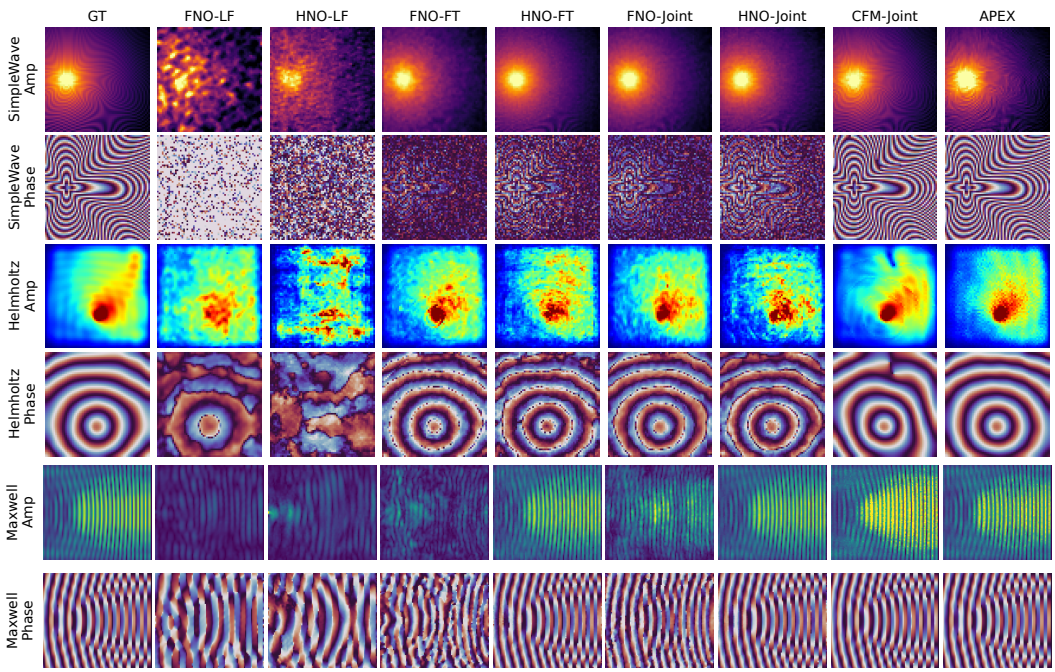


Figure 5: Additional qualitative comparison at HF50.

results for the three higher-frequency target groups, HF20, HF50, and HF100, on the SimpleWave, Helmholtz, and Maxwell benchmarks. For space reasons, the Overall summary column is omitted here. Lower $H^1$ and higher AWPC are better.

## G LIMITATIONS

A current limitation of APEX is its inference cost. Since the proposed enhancer is based on conditional flow matching, generating a higher-frequency field requires iterative sampling and is therefore slower than a single forward pass of a neural operator. In this work, we focus on improving prediction accuracy under scarce target-frequency supervision, and do not optimize the sampling efficiency in detail.

Another limitation is that we instantiate the lower-frequency backbone using FNO in the main APEX framework. The proposed amplitude-anchor design is not restricted to FNO in principle, since any lower-frequency surrogate that produces a coarse target-frequency prediction can be used to provide the amplitude anchor. A systematic evaluation of alternative neural-operator backbones and physics-aware surrogate architectures is left for future work.

Finally, the Green's-function-inspired phase prior used in this work is intentionally simplified. It provides source-aware and geometry-aware oscillatory cues, but does not fully model complex multipath effects, strong dispersion, or detailed boundary interactions. Future work may incorporate more accurate physics-based priors or data-calibrated phase models.

Table 7: Bootstrap confidence intervals for the higher-frequency prediction results across three benchmarks. Each target group reports mean ± half-width of a 95% non-parametric bootstrap confidence interval for $H^1$/AWPC over the fixed HF test split ($B$=2000). Underlined entries denote the best result within each dataset and metric column.

**SimpleWave**

| Method | HF20 | | HF50 | | HF100 | |
|---|---|---|---|---|---|---|
| | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ |
| FNO-LF | 1.377 ± 0.005 | 0.012 ± 0.001 | 1.454 ± 0.005 | 0.013 ± 0.001 | 1.429 ± 0.005 | 0.013 ± 0.001 |
| FNO-FT | 1.053 ± 0.010 | 0.157 ± 0.014 | 1.114 ± 0.018 | 0.238 ± 0.017 | 1.128 ± 0.011 | 0.133 ± 0.011 |
| HNO-LF | 1.359 ± 0.005 | 0.012 ± 0.001 | 1.420 ± 0.005 | 0.014 ± 0.001 | 1.401 ± 0.005 | 0.014 ± 0.001 |
| HNO-FT | 0.923 ± 0.032 | 0.496 ± 0.029 | 1.021 ± 0.033 | 0.465 ± 0.027 | 1.071 ± 0.031 | 0.400 ± 0.028 |
| FNO-Joint | 0.963 ± 0.026 | 0.403 ± 0.026 | 1.032 ± 0.033 | 0.438 ± 0.026 | 1.100 ± 0.027 | 0.324 ± 0.025 |
| HNO-Joint | 0.885 ± 0.038 | 0.540 ± 0.031 | 0.973 ± 0.039 | 0.531 ± 0.030 | 1.044 ± 0.034 | 0.433 ± 0.031 |
| CFM-Joint | 0.935 ± 0.048 | 0.620 ± 0.034 | 1.151 ± 0.046 | 0.516 ± 0.031 | 1.234 ± 0.030 | 0.351 ± 0.025 |
| APEX | 0.396 ± 0.004 | 0.909 ± 0.002 | 0.494 ± 0.013 | 0.865 ± 0.006 | 0.804 ± 0.020 | 0.659 ± 0.018 |

**Helmholtz**

| Method | HF20 | | HF50 | | HF100 | |
|---|---|---|---|---|---|---|
| | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ |
| FNO-LF | 1.611 ± 0.008 | 0.169 ± 0.005 | 1.686 ± 0.008 | 0.098 ± 0.004 | 2.261 ± 0.029 | 0.050 ± 0.002 |
| FNO-FT | 0.762 ± 0.005 | 0.897 ± 0.002 | 0.783 ± 0.005 | 0.874 ± 0.003 | 0.966 ± 0.009 | 0.759 ± 0.005 |
| HNO-LF | 2.989 ± 0.027 | 0.097 ± 0.004 | 3.743 ± 0.041 | 0.042 ± 0.002 | 5.637 ± 0.097 | 0.032 ± 0.002 |
| HNO-FT | 0.756 ± 0.005 | 0.871 ± 0.002 | 0.909 ± 0.007 | 0.821 ± 0.003 | 1.276 ± 0.021 | 0.696 ± 0.006 |
| FNO-Joint | 0.715 ± 0.006 | 0.903 ± 0.002 | 0.779 ± 0.007 | 0.850 ± 0.003 | 0.899 ± 0.010 | 0.737 ± 0.006 |
| HNO-Joint | 0.930 ± 0.008 | 0.853 ± 0.002 | 1.006 ± 0.010 | 0.794 ± 0.003 | 1.401 ± 0.025 | 0.662 ± 0.005 |
| CFM-Joint | 0.852 ± 0.006 | 0.653 ± 0.014 | 0.895 ± 0.006 | 0.543 ± 0.015 | 0.930 ± 0.007 | 0.397 ± 0.016 |
| APEX | 0.422 ± 0.010 | 0.928 ± 0.003 | 0.492 ± 0.013 | 0.885 ± 0.005 | 0.634 ± 0.019 | 0.815 ± 0.008 |

**Maxwell**

| Method | HF20 | | HF50 | | HF100 | |
|---|---|---|---|---|---|---|
| | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ | $H^1 \downarrow$ | AWPC↑ |
| FNO-LF | 1.090 ± 0.009 | 0.066 ± 0.011 | 1.023 ± 0.002 | 0.003 ± 0.004 | 1.008 ± 0.000 | 0.000 ± 0.002 |
| FNO-FT | 0.840 ± 0.017 | 0.465 ± 0.020 | 0.899 ± 0.014 | 0.301 ± 0.022 | 0.954 ± 0.007 | 0.169 ± 0.013 |
| HNO-LF | 1.097 ± 0.007 | 0.044 ± 0.012 | 1.045 ± 0.003 | 0.012 ± 0.006 | 1.014 ± 0.001 | 0.002 ± 0.002 |
| HNO-FT | 0.912 ± 0.026 | 0.421 ± 0.030 | 0.947 ± 0.026 | 0.329 ± 0.032 | 0.962 ± 0.021 | 0.277 ± 0.031 |
| FNO-Joint | 0.904 ± 0.020 | 0.422 ± 0.024 | 0.933 ± 0.020 | 0.337 ± 0.025 | 0.959 ± 0.012 | 0.233 ± 0.018 |
| HNO-Joint | 0.842 ± 0.025 | 0.441 ± 0.031 | 0.900 ± 0.026 | 0.342 ± 0.032 | 0.948 ± 0.021 | 0.285 ± 0.031 |
| CFM-Joint | 0.701 ± 0.029 | 0.828 ± 0.018 | 0.897 ± 0.033 | 0.680 ± 0.027 | 1.187 ± 0.034 | 0.411 ± 0.036 |
| APEX | 0.697 ± 0.034 | 0.812 ± 0.022 | 0.790 ± 0.031 | 0.725 ± 0.022 | 0.929 ± 0.029 | 0.624 ± 0.026 |