PASL-U: PHYSICS-ANCHORED OPERATOR LEARNING FOR SOURCE RECOVERY UNDER ENVIRONMENT SHIFTS

Anonymous authors
Paper under double-blind review

ABSTRACT

In guided-wave sensing, sensors record the source waveform that has been altered by propagation through the structure, so recovering the unknown source is an inverse problem that amounts to learning a mapping from measurements back to sources. However, the propagation operator changes with the environment: temperature changes the structure's stiffness, density, and wave speed, and thus changes the source-to-sensor propagation characteristics. Conventional operator learning cannot separate environmental effects from source effects, and therefore attributes the former to the source. Outside the calibration range, the influence of the environment on propagation gradually increases, causing source recovery to degrade. We propose PASL-U (Physics-Anchored Operator Learning for Source Recovery under Environment Shifts), which preserves the environmental dependence of a physical propagation prior, learns a bounded correction from calibration environments, and explicitly predicts how this correction changes with the environment. The corrected operator is then frozen and used for source recovery in unseen environments. Experiments show that PASL-U maintains lower propagation and source errors under strict scalar and coupled-parameter extrapolation within the tested range. On measured guided waves, it also retains damage-recognition performance in temperature bands outside the propagation-calibration range.

1 INTRODUCTION

Guided-wave structural health monitoring, which uses elastic waves guided along a thin structure and sensitive to local damage, infers unknown sources or damage from waves recorded by sparse sensors (Cawley & Alleyne, 1996; Mitra & Gopalakrishnan, 2016; Croxford et al., 2010; Moll et al., 2019). Learned reconstruction and algorithm unrolling provide data-driven approaches to this inference from measurements (Arridge et al., 2019; Monga et al., 2021). However, models trained for a particular physical setting may require new data and retraining when conditions change, limiting their reuse for inverse source recovery.

Guided waves are one instance of a broader problem class: source recovery under environment-dependent propagation in source-linear partial differential equation (PDE) systems. Operator learning addresses this need for reusable models by learning mappings across families of PDE inputs. DeepONet and Fourier neural operators predict solution fields for varying input functions and coefficients (Lu et al., 2021; Li et al., 2021). Environmental parameters can also enter these mappings, but sparse calibration leaves propagation at unseen conditions uncertain. For example, temperature or transport changes can modify source-to-sensor propagation, allowing environmental effects to obscure source-related changes (Croxford et al., 2010; Moll et al., 2019).

Physical knowledge helps constrain these mappings. Physics-informed neural networks and operators use governing-equation residuals to regularize learning (Raissi et al., 2019; Li et al., 2024). Neural Green's operators and learned forward–adjoint correction combine approximate physics with learned model corrections (Melchers et al., 2026; Lunz et al., 2021). These approaches constrain predictions through equations or compensate for discrepancies between approximate models and observed propagation. However, consistency with nominal physics does not eliminate model mismatch, and corrections inferred from a few calibrated environments may become inaccurate as en-

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.

vironmental parameters change. Residual propagation mismatch can then be misattributed to the source and amplified during inversion. The remaining challenge is to extend calibrated corrections across environments while accounting for their effect on source recovery. This motivates treating environmental extrapolation as a propagation-operator prediction problem rather than as direct extrapolation of the inverse source map.

We therefore study a specific operator-learning question: how to predict a calibrated correction to a misspecified physical operator beyond observed environments, while jointly accounting for its effect on source recovery. To address it, we propose PASL-U, a two-stage correction-and-recovery framework for source-linear PDE systems. Our contributions are:

1. We propose PASL-U, which represents the environment-dependent propagation operator as a physical prior plus a bounded, environment-conditioned correction. This allows source-linear propagation to be extrapolated beyond the calibration range while retaining the environmental dependence of the physical prior.
2. We explicitly separate correction-family approximation from environment-code extrapolation and derive a conditional error decomposition that identifies how these errors are amplified during operator inversion.
3. We evaluate strict extrapolation under scalar and coupled environmental shifts, together with synthetic and measured guided-wave data. The results show consistent reductions in transfer-operator and source-recovery error within the tested extrapolation regimes.

2 Related Work

Research relevant to PASL-U connects three themes: operator learning across PDE parameters, explicit propagation modeling and operator correction, and physics-informed inverse operator learning. These research directions address complementary aspects of prediction and reconstruction in physical systems. Our focus is their interaction when environmental changes alter propagation and the available physical model is imperfect.

Operator learning across PDE parameters. DeepONet and Fourier neural operators learn mappings between function spaces, allowing solution prediction across families of PDE inputs (Lu et al., 2021; Li et al., 2021). Physics-informed neural operators incorporate governing-equation constraints alongside data to improve solution approximation (Li et al., 2024). Extending prediction beyond the training distribution has also received explicit attention. Zhu et al. (2023) develop extrapolation methods based on fine-tuning and multifidelity learning, using PDE information or sparse new observations. PASL-U considers environmental extrapolation with a potentially misspecified physical prior. It preserves the prior's environmental dependence and predicts a calibrated discrepancy, with propagation-model parameters remaining fixed during deployment.

Explicit propagation modeling and operator correction. Green's representations expose the mapping from excitation to response. Praveen et al. (2023) construct low-rank Green's approximations from excitation–response data and interpolate their modes across model parameters. Neural Green's operators preserve linear dependence on source terms while learning nonlinear dependence on PDE coefficients; their physics-based constructions include learned corrections to truncated Neumann approximations of the inverse system matrix (Melchers et al., 2026). For inverse problems, Lunz et al. (2021) learn separate forward and adjoint corrections and train them along reconstruction trajectories. PASL-U places a bounded discrepancy in the discrete physical operator before inversion. Shared correction modes and calibration coordinates are identified first, followed by an environment-to-coordinate map learned with the modes frozen. The resulting source-linear transfer and its exact adjoint are then used throughout source recovery.

Physics-informed inverse operator learning. Physics-informed and data-driven approaches have also been applied directly to inverse problems. PINNs embed governing-equation residuals to solve forward and inverse problems in one optimization (Raissi et al., 2019). Physics-informed DeepONets learn solution operators for parametric PDEs under governing-equation constraints (Wang et al., 2021). Data-driven inverse solvers and algorithm unrolling learn reconstruction maps from paired measurements (Arridge et al., 2019; Monga et al., 2021). PI-DION learns inverse solution

operators from partial measurements through PDE constraints and observation consistency, without paired target labels (Cho & Son, 2025). These methods constrain or learn the inverse mapping with known physics, but they do not explicitly model how a misspecified propagation operator changes across environments.

3 PASL-U: Physics-Anchored Operator Extrapolation

We propose PASL-U, a unified framework for inverse source recovery under environment-dependent propagation, as illustrated in Figure 1. The central idea is to retain the environmental dependence of an evaluable physical prior and learn a bounded correction to its operator before inversion. Thus, PASL-U does not require the learned component to represent the full propagation operator across environments; it only needs to represent the discrepancy left by the physical prior. We parameterize this correction using a shared offset and correction modes, weighted by coefficients that vary with the environment. The corrected operator then defines a source-to-sensor map that preserves linearity in the source while incorporating the discrepancy learned from calibration data.

Learning proceeds in two stages. Stage A jointly estimates the shared correction components and coefficients at calibration environments from known probes, using propagation objectives. Stage B freezes these components and learns a bounded coefficient predictor that combines a linear environmental trend with a nonlinear residual. Stage A therefore defines the correction family that can be represented, whereas Stage B determines how this family is indexed by the environment. The two stages are separated to expose these two error sources explicitly; the separation is not assumed to be necessary for predictive accuracy. This predictor is trained against the Stage-A coefficients while retaining calibration consistency and physical-admissibility constraints. At a new environment, the physical prior and predicted correction determine the propagation map, which remains fixed during source recovery. Proximal-gradient iterations use this map and its exact adjoint to recover group-sparse source waveforms. Different PDE systems provide their own physical realizations and structural constraints and are calibrated separately within the same framework. We now formalize this framework, beginning with the problem setting and the operators on which it builds.

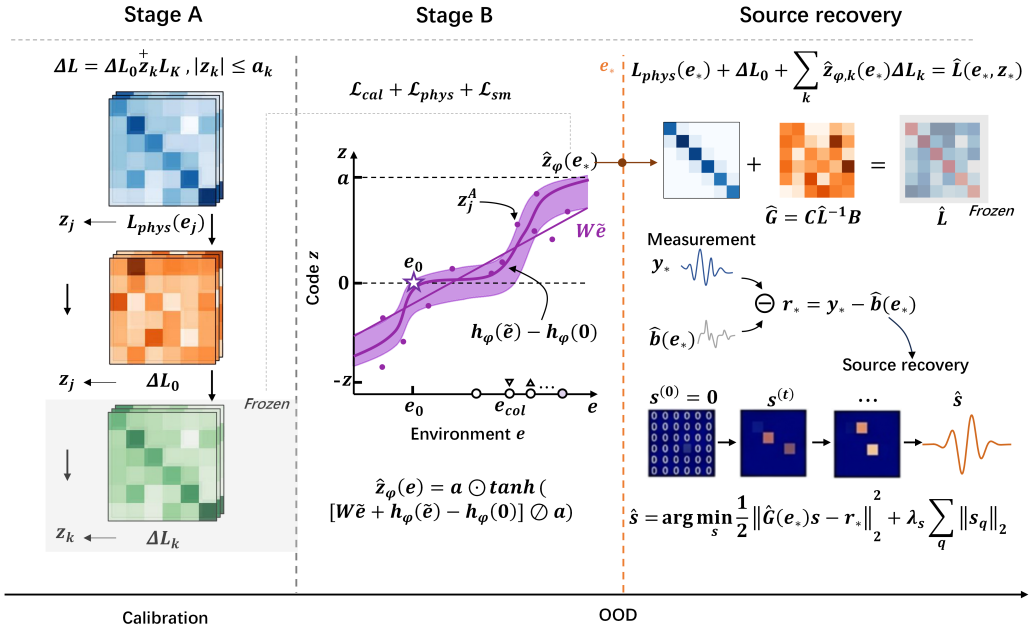


Figure 1: Overview of PASL-U for physics-anchored operator extrapolation and source recovery. Stage A learns a shared bounded correction family, and Stage B learns its environment-to-code mapping. At unseen environments, the predicted correction is combined with the physical prior to construct the propagation operator for source recovery.

3.1 PROBLEM SETTING

Let $e \in \mathbb{R}^{d_e}$ contain d_e observed environmental variables, such as temperature or transport parameters. With Q candidate source locations, P measurement channels, and T time samples, a source $s \in \mathbb{R}^{Q\times T}$ produces measurements $y \in \mathbb{R}^{P\times T}$ through

$$y = \mathcal{G}^\star(e)s + b(e) + \varepsilon, \tag{1}$$

where $\mathcal{G}^\star(e)$ is the true, environment-dependent propagation operator, $b(e)$ is an optional structured background (such as the intact-structure response in guided waves), and ε is measurement noise. We are given calibration data $\mathcal{D}_{\text{cal}} = \{(e_j, x_{jm}, y_{jm})\}$, where j indexes the calibration environment and m indexes probe records within it: x_{jm} is a known probe waveform and y_{jm} is its measured response. The objective is to recover an unknown source from its measurement and observed environment $e_\star$, both within and beyond the calibration range, without additional probes or updates to the propagation model.

We assume propagation is linear in the source and that an approximate physical operator $L_{\text{phys}}(e,\omega) \in \mathbb{C}^{N\times N}$ can be evaluated at calibration and test environments, where ω is frequency and N is the state dimension. Given a source-injection operator $B \in \mathbb{C}^{N\times Q}$ and a sensor-extraction operator $C \in \mathbb{C}^{P\times N}$, the approximate frequency-domain transfer is $G_{\text{phys}} = C\,L_{\text{phys}}^{-1}B$. We write $G^\star(e,\omega)$ for the true frequency-domain transfer, so that $\mathcal{G}^\star(e)$ is its time-domain action. The physical operator may contain structured model discrepancy (Kennedy & O'Hagan, 2001), which we model over the range of interest with bounded corrections parameterized by low-dimensional environment codes. From here on we suppress the frequency argument: G denotes a transfer matrix and $\mathcal{G}$ its time-domain action.

We use $\odot$ and $\oslash$ for entrywise multiplication and division, X^H for conjugate transpose, and A^* for the Euclidean adjoint. The norm $\|\cdot\|_2$ is Euclidean for signals (after vectorizing waveform arrays) and spectral for operators; $\|\cdot\|_F$ is the Frobenius norm.

3.2 PHYSICS-ANCHORED CORRECTION FAMILY

To correct the misspecified physical prior while preserving source linearity and the prior's environment dependence, this section maps an environment code z to a corrected operator $\widehat{L}(e,z)$ and its transfer $\widehat{G}(e,z)$. We parameterize the correction before forming the transfer matrix:

$$\begin{aligned} &\widehat{L}(e,z) = L_{\text{phys}}(e) + \Delta L(z), \qquad \Delta L(z) = \Delta L_0 + \sum_{k=1}^{K} z_k \Delta L_k, \\ &\widehat{G}(e,z) = C\widehat{L}(e,z)^{-1}B. \end{aligned} \tag{2}$$

Here ΔL_0 is a shared bias, $\{\Delta L_k\}_{k=1}^K$ are learned modes, and $z \in \mathbb{R}^K$ specifies their environment-dependent combination. K controls family dimension; individual modes need not have low matrix rank. Fixed B and C retain the supplied geometry. Matrix inversion converts this affine correction into a generally nonlinear transfer family.

To control departure from the physical prior, we restrict codes to $\mathcal{Z} = \{z : |z_k| \le a_k\}$, with $a = (a_1,\dots,a_K)^\top \in \mathbb{R}^K_{>0}$. A mask $M \in [0,1]^{N\times N}$ specifies allowed corrections, $S \in \mathbb{R}^{N\times N}_{>0}$ supplies entrywise physical scales, and $\rho > 0$ sets the correction budget. The code-independent normalization in Appendix A.1 preserves shared modes and guarantees the entrywise bound

$$|\Delta L(z)| \le \rho M \odot S, \qquad z \in \mathcal{Z}. \tag{3}$$

The budget controls the learned discrepancy at every environment, including extrapolation, while preserving the environmental variation supplied by $L_{\text{phys}}(e)$. Conditions for controlling the inverse are given below.

3.3 LEARNING THE CORRECTION AND ITS PARAMETER DEPENDENCE

Stage A: calibrating a shared operator subspace. Stage A jointly learns shared modes and a free code for each calibration environment, fitting the variations supported by measurements before

modeling their dependence on e. We use $z_j = a \odot \tanh(r_j - r_{j_0})$, where $r_j \in \mathbb{R}^K$ are free variables and j_0 indexes the calibration environment nearest the reference $e_0 \in \mathbb{R}^{d_e}$. Setting $z_{j_0} = 0$ assigns the reference correction to ΔL_0, fixing a translation ambiguity without making the modes or physical coefficients uniquely identifiable.

Source-conditioned probe banks pair known spatial input waveforms with their source-induced responses and yield regularized transfer snapshots. Direct-action records constrain each measurement path through its available reference input. Calibration combines a transfer-matching loss $\mathcal{L}_{\mathrm{tr}}$ and a response-matching loss $\mathcal{L}_{\mathrm{act}}$, weighted by λ_{tr} and λ_{act}, over applicable records:

$$\mathcal{L}_{\mathrm{cal}} = \lambda_{\mathrm{tr}}\mathcal{L}_{\mathrm{tr}} + \lambda_{\mathrm{act}}\mathcal{L}_{\mathrm{act}}. \tag{4}$$

Stage A minimizes this calibration loss together with penalties on correction magnitude and violations of the supplied physical structure. Appendix A.2 gives the full objective and Appendix A.3 specifies the physical penalties.

Stage B: learning an environment-to-code map. To extrapolate the correction to unseen environments, this section learns a map from the environment e to the predicted code $\widehat{z}_\phi(e)$. We freeze the modes and fit a map to the calibrated codes $\{z_j^A\}$. With normalized environment $\widetilde{e} = (e - e_0) \oslash s_e$, the map combines a linear path with an anchored MLP residual:

$$\widehat{z}_\phi(e) = a \odot \tanh([W\widetilde{e} + h_\phi(\widetilde{e}) - h_\phi(0)] \oslash a)\,. \tag{5}$$

Here $W \in \mathbb{R}^{K \times d_e}$, $h_\phi : \mathbb{R}^{d_e} \to \mathbb{R}^K$, and $s_e \in \mathbb{R}^{d_e}_{>0}$. The parameter set ϕ includes both W and the MLP parameters. The linear path represents a local trend, and the MLP residual allows nonlinear dependence. The anchor gives $\widehat{z}_\phi(e_0) = 0$, coinciding with Stage A when e_0 is a calibration environment. The final $\tanh$ preserves the code box, although saturation limits the representation of a growing discrepancy.

Stage B optimizes ϕ by matching Stage-A codes and maintaining calibration consistency through the frozen modes. Physical penalties also query prespecified unlabeled environments, including a margin outside calibration, and a smoothness term regularizes the predicted codes. These queries use the physical prior without test measurements or target transfers. The complete objective and finite-difference penalty are given in Appendix A.2 and Eq. 17. Supplied validation data are used for selection.

3.4 Extrapolation and source recovery

To recover the source at an unseen environment, this section assembles the frozen transfer $\widehat{G}(e_\star)$ and solves a group-sparse inverse problem for the estimated source $\widehat{s}$. At deployment, substituting $\widehat{z}_\phi(e_\star)$ into Eq. 2 gives the frozen transfer $\widehat{G}(e_\star) = \widehat{G}(e_\star, \widehat{z}_\phi(e_\star))$. For a background-subtracted observation $r_\star = y_\star - \widehat{b}(e_\star)$, where $\widehat{b}(e_\star)$ is the predicted background, source recovery solves

$$\widehat{s} = \arg\min_s \frac{1}{2}\|\widehat{\mathcal{G}}(e_\star)s - r_\star\|_2^2 + \lambda_s \sum_{q=1}^{Q} \|s_q\|_2, \tag{6}$$

where $\widehat{\mathcal{G}}(e_\star)$ is the time-domain action of $\widehat{G}(e_\star)$, s_q is the waveform at location q, and $\lambda_s \geq 0$. Positive λ_s encourages sparsity across source locations through the group-lasso penalty (Yuan & Lin, 2006); $\lambda_s = 0$ gives least squares. Proximal-gradient updates use the fixed predicted operator and its exact adjoint (Parikh & Boyd, 2014).

Synthetic guided waves calibrate source transfer and intact background separately, then profile a scalar background amplitude (Appendix B.3). Measured guided waves apply intact-calibrated path corrections to both the physical candidate-source transfer and the direct response. The corrected direct response to each record's pitch-derived input predicts the background to subtract. Candidate geometry and solver settings are specified in Appendix C.2.

3.5 Theoretical analysis

We now explain why the two-stage structure remains stable when the environment leaves the calibration range. The analysis follows one decomposition chain: source-recovery error is governed by

transfer-operator error, transfer-operator error by the correction-family and environment-code errors, and the code error grows linearly with environmental distance. Full proofs are in Appendix B.

Fix the learned modes and an environment e, and assume invertibility on the code box, $\sup_{z\in\mathcal{Z}} \|\widehat{L}(e,z)^{-1}\|_2 \le \kappa_e < \infty$. Let $z^\dagger(e) \in \mathcal{Z}$ be an ideal comparison code, and let $\epsilon_{\mathrm{fam}}(e) = \|\widehat{G}(e, z^\dagger(e)) - G^\star(e)\|_2$ be the correction-family approximation error. Since the code-to-transfer map $z \mapsto \widehat{G}(e,z)$ is Lipschitz with constant $c_e = \|C\|_2\|B\|_2\kappa_e^2(\sum_k \|\Delta L_k\|_2^2)^{1/2}$ (Appendix B.1),

$$\left\|\widehat{G}(e) - G^\star(e)\right\|_2 \le \epsilon_{\mathrm{fam}}(e) + c_e \left\|\widehat{z}_\phi(e) - z^\dagger(e)\right\|_2 . \tag{7}$$

The first term is the best approximation error of the correction family, reduced by Stage-A's modes; the second is Stage-B's code-prediction error, amplified by inverse conditioning through c_e. The two stages thus control the two error sources in the same bound. A budget condition $\theta < 1$, with $\theta = p\beta$, $p = \|L_{\mathrm{phys}}(e)^{-1}\|_2$, and $\beta = \rho\|M \odot S\|_F$, bounds κ_e (Appendix B.1).

Assume the predictor $\widehat{z}_\phi$ is L_ϕ-Lipschitz and the comparison map $z^\dagger$ is $L_\dagger$-Lipschitz on a convex domain containing e_b and e. Then

$$\left\|\widehat{z}_\phi(e) - z^\dagger(e)\right\|_2 \le \left\|\widehat{z}_\phi(e_b) - z^\dagger(e_b)\right\|_2 + (L_\phi + L_\dagger)\,\|e - e_b\|_2, \tag{8}$$

which, substituted into Eq. 7, gives the distance-aware transfer bound

$$\left\|\widehat{G}(e) - G^\star(e)\right\|_2 \le \epsilon_{\mathrm{fam}}(e) + c_e \left\|\widehat{z}_\phi(e_b) - z^\dagger(e_b)\right\|_2 + c_e(L_\phi + L_\dagger)\,\|e - e_b\|_2. \tag{9}$$

The right-hand side separates family-representation error, reference code error, and a term growing linearly with environmental distance, so the extrapolation error grows linearly rather than abruptly (Appendix Figure 3). The reference code error admits a finite-sample bound (Appendix B.2), and the transfer error propagates to source recovery through a Parseval identity (Appendix B.3).

4 Experiments

We empirically evaluate PASL-U's ability to model environment-dependent propagation and recover sources on three inverse problems: strict out-of-distribution (OOD) advection–diffusion source recovery, dual-parameter advection–diffusion source recovery, and synthetic guided-wave source recovery. We further validate its practical application on measured OGW2 guided-wave damage recognition. For comparison, we use Physical-only, Data-NGO, Neumann NGO, Lunz-FAC, PINO-style, DeepONet, and FNO as baselines. Training details, including data splits, training budgets, solver settings, and the calibration–validation protocol, are given in Appendix C. Additional experiments, including core-module ablation, are reported in Appendix C.3. Across all four experiments, PASL-U obtains richer propagation-calibration information under known environmental parameters, thereby maintaining stable source recovery, and preserves more information useful for recovery and recognition in unseen environments, leading to more robust extrapolated recovery and damage recognition.

4.1 Strict OOD advection–diffusion source recovery

We first consider the transient advection–diffusion equation on a two-dimensional square domain $\Omega = [-1,1]^2$,

$$u_t - \nabla\cdot\big(\kappa(x,\mu)\nabla u\big) + v(\mu)\cdot\nabla u + \rho(x,\mu)u = s(x,t),$$

where $u(x,t)$ is the concentration or pressure perturbation to be recovered and $s(x,t)$ is the source term. Zero Dirichlet conditions are imposed on the outer boundary and the initial state is zero. The environmental parameter μ enters the diffusivity $\kappa(x,\mu)$, advection velocity $v(\mu)$, and reaction coefficient $\rho(x,\mu)$, so it changes source-to-sensor propagation rather than the spatial placement of sources or sensors. We discretize the spatial operator with five-point finite differences and the temporal derivative with backward Euler, obtaining a frequency-domain source-to-sensor transfer operator $P(\mu)$. The physical prior uses known μ-dependent coefficient forms, while the true propagation model additionally contains μ-dependent spatial modes; the physical prior is therefore approximate but not exact.

In this experiment, a 21×21 spatial grid is observed by 64 interior sensors, with 25 candidate source locations. Calibration and validation use parameters in $[0.6, 1.2]$; we evaluate strict extrapolation at $\mu \in \{1.4, 1.6, 1.8, 2.0\}$, disjoint from the calibration range. Thus, all reported OOD environments require extrapolation beyond the largest calibrated parameter. Table 1 reports two-source recovery means across the four OOD environments, without standard deviations; per-parameter single-source results are reported in Appendix C.4. Figure 2 shows the relative transfer-operator error of the source-linear methods outside the calibration range.

FNO and DeepONet exhibit substantially larger source NMSE under the evaluated strict-OOD protocol, remaining close to 1 across the four extrapolation environments. This behavior is consistent with their direct inverse formulation, which does not explicitly parameterize the environment-dependent source-to-sensor propagation operator.

As the test environment continues to extrapolate, PASL-U attains the lowest relative transfer-operator error at every extrapolation μ and the lowest two-source NMSE at the farther environments ($\mu = 1.8, 2.0$), while Neumann NGO is competitive at the nearer environments ($\mu = 1.4, 1.6$). This indicates that PASL-U, by learning a propagation operator that varies with the environmental parameter, retains more source information as extrapolation distance grows and therefore achieves the most stable source recovery far from calibration.

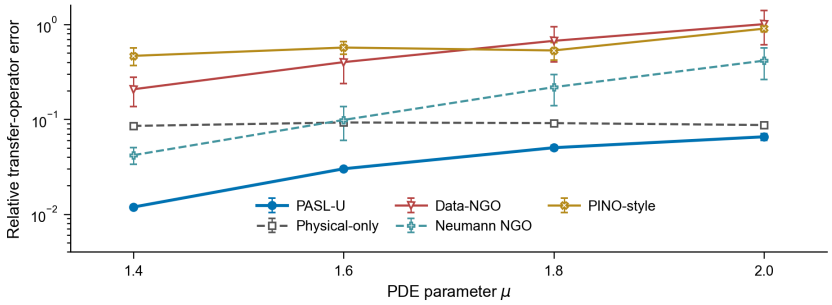


Figure 2: Relative transfer-operator error for the five source-linear operator methods under strict OOD extrapolation (lower is better).

Table 1: Advection–diffusion two-source NMSE at 30 dB under strict OOD extrapolation, five-seed mean (lower is better).

Method	$\mu = 1.4$	$\mu = 1.6$	$\mu = 1.8$	$\mu = 2.0$
Physical-only	0.5142	0.5056	0.4898	0.4683
PASL-U	0.4760	0.4544	**0.4309**	**0.4059**
Data-NGO	0.5277	0.5591	0.6091	0.6528
Neumann NGO	**0.4745**	**0.4485**	0.4355	0.4800
Lunz-FAC	0.5159	0.5078	0.4925	0.4717
PINO-style	0.6500	0.6717	0.6450	0.7192
DeepONet	0.8829	0.8993	0.9078	0.9027
FNO	1.0010	1.0055	1.0119	1.0151

4.2 Dual-parameter extrapolation

To test whether the observed extrapolation extends beyond a single environmental coordinate, we introduce a coupled two-parameter transport setting in which c and d jointly change diffusivity, advection, and reaction. Parameter c controls the base propagation fields:

$$P_c(c) = C\,\mathcal{L}\big[\kappa^\star(x,c),\, v^\star(c),\, \rho^\star(x,c)\big]^{-1} B.$$

Parameter d then acts on those fields as

$$P(c,d) = C\,\mathcal{L}\Big[\kappa^\star(x,c)\,(1+\alpha_\kappa d + \beta_\kappa c d),\; v^\star(c) + \gamma_d\, d,\; \rho^\star(x,c)\,(1+\alpha_\rho d)\Big]^{-1} B,$$

where the true base fields $\kappa^\star(x, c)$, $v^\star(c)$, and $\rho^\star(x, c)$ contain c-dependent hidden spatial modes. Calibration uses $c, d \in \{0.8, 0.95, 1.1, 1.2\}$ (16 environments), and evaluation uses eight extrapolation environments up to $[2.0, 2.0]$. Training uses 2000 epochs with early stopping, and recovery uses group-sparse reconstruction with regularization 0.05 and 150 iterations. Values are five-seed mean $\pm$ sample SD.

Table 2: Two-source recovery NMSE by extrapolation environment at 30 dB, five-seed mean (lower is better).

Environment	PASL-U	Lunz-FAC	Physical-only	PINO-style	Data-NGO	FNO	DeepONet	Neumann-NGO
$[1.6, 1.0]$	**0.2699**	0.2907	0.3129	0.3551	0.5680	0.7528	0.7553	0.3782
$[2.0, 1.0]$	**0.2671**	0.2884	0.3087	0.3803	0.6538	0.7519	0.7474	2.5437
$[1.0, 1.6]$	**0.2792**	0.3090	0.3089	0.3421	0.5375	0.7620	0.7645	0.3022
$[1.0, 2.0]$	**0.2719**	0.3143	0.3120	0.3640	0.5867	0.7449	0.7525	0.2999
$[1.4, 1.4]$	0.2880	0.3128	0.3327	0.3476	0.5674	0.7561	0.7573	**0.2619**
$[1.6, 1.6]$	**0.2748**	0.3036	0.3239	0.4336	0.6225	0.7401	0.7561	0.5010
$[1.8, 1.8]$	**0.2778**	0.3065	0.3268	0.5289	0.6887	0.7590	0.7575	2.0811
$[2.0, 2.0]$	**0.2778**	0.3041	0.3242	0.4714	0.7382	0.7607	0.7583	2.8316

PASL-U maintains a narrow range of two-source NMSE across the eight extrapolation environments, with values between 0.267 and 0.288, including the farthest $[2.0, 2.0]$ environment, as shown in Table 2. Neumann-NGO collapses at the far end, reaching 2.54 at $[2.0, 1.0]$ and 2.83 at $[2.0, 2.0]$. FNO and DeepONet remain near 0.75 everywhere and cannot extrapolate. The propagation-operator relative error is reported in Appendix C.5.

The two extrapolation directions are not equivalent. Parameter c controls the spatially varying diffusivity, velocity, and reaction terms inherited from the single-parameter system, while d enters primarily as a smooth multiplicative scaling and a small linear velocity offset. Extrapolating along c therefore changes the spatial structure of the true transfer operator more strongly and exposes hidden spatial modes that the physical prior misses. Extrapolating along d remains closer to a smooth overall scaling, for which the physical prior and low-rank physical corrections are more adequate. This asymmetry explains why Neumann-NGO remains competitive at $[1.0, 1.6]$ and $[1.0, 2.0]$, where d dominates, but collapses at $[2.0, 1.0]$ and $[2.0, 2.0]$, where c increases. In contrast, PASL-U explicitly predicts an environment-dependent correction code for both parameters, allowing it to retain the physical dependence in each direction and keep two-source NMSE within 0.267–0.288 across all environments.

4.3 Measured guided waves: damage recognition

To explore whether PASL-U can retain source-recovery ability under parameter extrapolation on real data without source ground truth, we evaluate it on measured OGW2 guided-wave damage recognition. Before the formal experiments, we first simulate the setting on synthetic data, as described in Appendix C.6. The measured pipeline calibrates propagation using intact records below $45°$C. The frozen inverse produces four 256-sample source traces, which are flattened into 1024 features for a linear classifier. Damage labels below the same temperature boundary train and select the classifier, while higher-temperature records are used for evaluation. Table 3 summarizes the damage-recognition accuracy in each temperature band.

Table 3: Measured OGW2 damage-recognition accuracy, five-seed mean $\pm$ sample SD.

Method	45–50°C	50–55°C	55–60°C
Physical-only	53.8 ± 2.3	25.0 ± 0.0	25.0 ± 0.0
PASL-U	$\mathbf{86.0 \pm 20.4}$	$\mathbf{57.0 \pm 28.8}$	$\mathbf{33.8 \pm 12.1}$
Data-NGO	54.5 ± 3.0	25.0 ± 0.0	25.0 ± 0.0
Neumann NGO	60.5 ± 10.0	35.8 ± 6.2	27.1 ± 2.7
PINO-style	29.2 ± 4.5	13.8 ± 7.7	16.7 ± 8.3
DeepONet	25.0 ± 0.0	25.0 ± 0.0	25.0 ± 0.0
FNO	46.8 ± 3.3	25.0 ± 0.0	25.0 ± 0.0
Lunz-FAC	43.0 ± 10.0	25.0 ± 0.0	25.0 ± 0.0

From Table 3, PASL-U attains the highest accuracy in all three temperature bands, and its accuracy degrades clearly more slowly than every baseline. This advantage is most pronounced in the middle extrapolation band: at 50–55°C, PASL-U retains 57.0% recognition accuracy, whereas most baselines have already degraded to near the four-class random level. The direct inverse baselines FNO and DeepONet remain near chance across all bands, indicating substantially degraded extrapolation performance under the evaluated distribution shift—as expected, because they directly map high-dimensional observations to sources without explicitly modeling the environment-dependent propagation operator and therefore cannot exploit the underlying PDE structure for extrapolation.

Notably, PASL-U's advantage stems from its two-stage environment encoding: Stage A learns bounded correction modes at calibration environments, and Stage B explicitly learns an environment-to-code predictor. This allows it to predict how the propagation operator changes once the environment leaves the calibration range, thereby retaining more damage-related source information. By contrast, methods such as Data-NGO rely primarily on flexible fitting within the calibration range and lose support once extrapolated, while the physical-prior methods have a pathwise effective prior with limited physical constraints, so the recovered features lose discriminability as the extrapolation distance grows.

5 Limitations and Discussion

Across the experiments and the error bound, the same mechanism explains the extrapolation behavior: what limits recovery far from calibration is the environment-to-code prediction, not the capacity of the correction family. Modeling how the propagation operator changes with the environment therefore matters more than making the correction family more flexible.

Physical prior and conditioning. Extrapolation depends on the physical prior, the correction budget, and inverse conditioning. We do not prove training convergence, and recovery is not covered when the budget fails to keep the operator invertible or the operator is ill-conditioned.

Environment coverage. The experiments cover scalar and two-parameter environments over finite ranges with multiple fitting runs. Jointly varying parameters, higher-dimensional parameter spaces, and independent measurement campaigns have not been tested.

Measured data. Measured guided waves calibrate propagation from intact records and evaluate damage recognition at temperatures outside the propagation-calibration range. Source-waveform ground truth and collocation beyond the prior-only placement are not covered.

6 Conclusion

PASL-U provides a unified physical-operator correction and source-inversion framework for source-linear PDE systems. On strict OOD transport extrapolation, operator correction reduces both propagation error and average source NMSE. In the dual-parameter transport setting, PASL-U keeps the lowest and most stable source recovery across all extrapolation environments, while Neumann-NGO collapses at the far end and the direct-inverse baselines cannot extrapolate. Synthetic guided-wave recovery results show waveform recovery at the nearest tested extrapolation temperature followed by weaker localization farther away. Measured recovery features further support damage recognition in temperature bands outside the propagation-calibration range. The observed extrapolation remains bounded by correction capacity, environment-code prediction, and operator conditioning, as discussed in the limitations.

REPRODUCIBILITY STATEMENT

Appendices A–C provide implementation details, proofs, data splits, evaluation protocols, and the archived evidence underlying each comparison. Anonymous code, experiment configurations, selected evidence, and verification scripts are provided in the supplementary material. Synthetic datasets are regenerated from the included manifests, while the measured OGW2 recordings are cited from (Moll et al., 2019) and are not redistributed. Run summaries and integrity hashes are included in the supplementary artifact; analyses whose protocols are not matched are explicitly separated from the scientific comparisons.

AI USE STATEMENT

Generative AI tools assisted manuscript revision, mathematical and implementation consistency checks, experiment and diagnostic code, and figure generation. Archived temperature- and noise-resolved plots are distinguished from newly executed training and recovery runs through separate experiment records. Numerical claims are based on recorded outputs; generated text or illustrative values are not used as experimental evidence.

REFERENCES

Simon Arridge, Peter Maass, Ozan Öktem, and Carola-Bibiane Schönlieb. Solving inverse problems using data-driven models. *Acta Numerica*, 28:1–174, 2019. doi: 10.1017/S0962492919000059.

Peter Cawley and David N. Alleyne. The use of Lamb waves for the long range inspection of large structures. *Ultrasonics*, 34(2–5):287–290, 1996. doi: 10.1016/0041-624X(96)00024-8.

Sung Woong Cho and Hwijae Son. Physics-informed deep inverse operator networks for solving PDE inverse problems. In *International Conference on Learning Representations*, 2025.

Anthony J. Croxford, Jochen Moll, Paul D. Wilcox, and Jennifer E. Michaels. Efficient temperature compensation strategies for guided wave structural health monitoring. *Ultrasonics*, 50(4–5):517–528, 2010. doi: 10.1016/j.ultras.2009.11.002.

Nicholas J. Higham. *Accuracy and Stability of Numerical Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2 edition, 2002. doi: 10.1137/1.9780898718027.

Marc C. Kennedy and Anthony O'Hagan. Bayesian calibration of computer models. *Journal of the Royal Statistical Society: Series B*, 63(3):425–464, 2001. doi: 10.1111/1467-9868.00294.

Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. In *International Conference on Learning Representations*, 2021. arXiv:2010.08895.

Zongyi Li, Hongkai Zheng, Nikola Kovachki, David Jin, Haoxuan Chen, Burigede Liu, Kamyar Azizzadenesheli, and Anima Anandkumar. Physics-informed neural operator for learning partial differential equations. *ACM / IMS Journal of Data Science*, 1(3):1–27, 2024. doi: 10.1145/3648506.

Karim Lounici, Massimiliano Pontil, Alexandre B. Tsybakov, and Sara van de Geer. Oracle inequalities and optimal inference under group sparsity. *The Annals of Statistics*, 39(4):2164–2204, 2011. doi: 10.1214/11-AOS896.

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. doi: 10.1038/s42256-021-00302-5.

Sebastian Lunz, Andreas Hauptmann, Tanja Tarvainen, Carola-Bibiane Schönlieb, and Simon Arridge. On learned operator correction in inverse problems. *SIAM Journal on Imaging Sciences*, 14(1):92–127, 2021. doi: 10.1137/20M1338460.

H. A. Melchers, J. H. M. Prins, and M. R. A. Abdelmalik. Neural Green's operators for parametric partial differential equations. *Computer Methods in Applied Mechanics and Engineering*, 455: 118893, 2026. doi: 10.1016/j.cma.2026.118893.

Mira Mitra and Srinivasan Gopalakrishnan. Guided wave based structural health monitoring: A review. *Smart Materials and Structures*, 25(5):053001, 2016. doi: 10.1088/0964-1726/25/5/053001.

Jochen Moll, Christian Kexel, Serena Pötzsch, Marcel Rennoch, and Axel S. Herrmann. Temperature affected guided wave propagation in a composite plate complementing the Open Guided Waves platform. *Scientific Data*, 6:191, 2019. doi: 10.1038/s41597-019-0208-1.

Vishal Monga, Yuelong Li, and Yonina C. Eldar. Algorithm unrolling: Interpretable, efficient deep learning for signal and image processing. *IEEE Signal Processing Magazine*, 38(2):18–44, 2021. doi: 10.1109/MSP.2020.3016905.

Neal Parikh and Stephen Boyd. Proximal algorithms. *Foundations and Trends in Optimization*, 1 (3):127–239, 2014. doi: 10.1561/2400000003.

Harshwardhan Praveen, Nicolas Boullé, and Christopher J. Earls. Principled interpolation of Green's functions learned from data. *Computer Methods in Applied Mechanics and Engineering*, 409: 115971, 2023. doi: 10.1016/j.cma.2023.115971.

Maziar Raissi, Paris Perdikaris, and George Em Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. doi: 10.1016/j.jcp.2018.10.045.

Sifan Wang, Hanwen Wang, and Paris Perdikaris. Learning the solution operator of parametric partial differential equations with physics-informed DeepONets. *Science Advances*, 7(40):eabi8605, 2021. doi: 10.1126/sciadv.abi8605.

Ming Yuan and Yi Lin. Model selection and estimation in regression with grouped variables. *Journal of the Royal Statistical Society: Series B*, 68(1):49–67, 2006. doi: 10.1111/j.1467-9868.2005.00532.x.

Min Zhu, Handi Zhang, Anran Jiao, George Em Karniadakis, and Lu Lu. Reliable extrapolation of deep neural operators informed by physics or sparse observations. *Computer Methods in Applied Mechanics and Engineering*, 412:116064, 2023. doi: 10.1016/j.cma.2023.116064.

A Implementation Details

A.1 Bounded correction parameterization

For bounded raw components $U_0, \dots, U_K$ and code bounds $a_k > 0$, the shared normalization is

$$D = \max\left(1, |U_0| + \sum_{k=1}^{K} a_k |U_k|\right), \qquad \Delta L_k = \rho M \odot S \odot U_k \oslash D, \quad k = 0, \dots, K. \tag{10}$$

All operations are entrywise. Since D is independent of z, the correction remains affine in the shared modes throughout $\mathcal{Z}$. The triangle inequality yields Eq. 3.

The raw components in Eq. 10 are mapped through $\tanh$ for real entries. For a complex entry u, the implementation uses $u \tanh(|u|)/\max(|u|, 10^{-8})$, yielding zero at $u = 0$. Both maps have entrywise magnitude at most one. Diagonal channel realizations store the diagonal entries only, while the advection–diffusion realization uses real matrices shared across frequencies. The normalization is applied to the stored entries and is independent of the environment code.

The Stage-A reference subtraction fixes $z_{j_0} = 0$ before decoding, so all calibration codes remain in $\mathcal{Z}$. Stage B is anchored at the configured e_0 through Eq. 5; the two reference locations coincide when e_0 is a calibration environment. An anchor fixes the reference correction, but the factorization into modes and codes can still be nonunique. Learning and evaluation concern the induced transfer operator.

A.2 Probe snapshots and calibration losses

For a probe bank, let M_j be the number of known probes at calibration environment j. The matrices $X_{j,f} \in \mathbb{C}^{Q \times M_j}$ and $Y_{j,f} \in \mathbb{C}^{P \times M_j}$ contain their inputs and measured responses at frequency f. Their regularized transfer snapshot is

$$G_{j,f}^{\mathrm{snap}} = Y_{j,f} X_{j,f}^H \left(X_{j,f} X_{j,f}^H + \gamma_f I\right)^{-1}, \tag{11}$$

where $\gamma_f > 0$ scales with probe energy. In the noiseless, unregularized case, a full-row-rank probe matrix identifies the transfer through a right inverse. With noise or incomplete excitation, regularization controls the estimate but does not establish full operator identifiability. Rank, conditioning, and reconstruction diagnostics accompany the snapshot. Direct-action records supply the measured response to their reference input. All calibration losses below are averaged over their applicable records.

For transfer targets, let W_j contain nonnegative entry weights, with the same shape as the transfer, and abbreviate the predicted calibration transfer by $\widehat{G}_j$. Here $W_j^{1/2}$ is the entrywise square root. The transfer loss is

$$\mathcal{L}_{\mathrm{tr}} = \frac{1}{N_{\mathrm{tr}}} \sum_j \frac{\|W_j^{1/2} \odot (\widehat{G}_j - G_j^{\mathrm{snap}})\|_F^2}{\max(\|W_j^{1/2} \odot G_j^{\mathrm{snap}}\|_F^2, \epsilon)}, \tag{12}$$

where N_{tr} is the number of available transfer records. For N_{act} probe-action records, the response loss is

$$\mathcal{L}_{\mathrm{act}} = \frac{1}{N_{\mathrm{act}}} \sum_{(e,x,y)} \frac{\|\widehat{\mathcal{G}}(e)x - y\|_2^2}{\max(\|y\|_2^2, \epsilon)}. \tag{13}$$

Training objectives. Stage A jointly optimizes the correction family and calibration codes using

$$\mathcal{L}_A = \mathcal{L}_{\mathrm{cal}} + \lambda_{\mathrm{prior}} \mathcal{L}_{\mathrm{prior}} + \lambda_{\mathrm{phys}}^A \mathcal{L}_{\mathrm{phys}}. \tag{14}$$

The prior penalty is the mean squared magnitude of $\Delta L(z_j) \oslash S$ over calibration codes and operator entries. The physical penalties are defined in Appendix A.3. Stage B freezes all Stage-A parameters and optimizes only ϕ:

$$\begin{aligned} \mathcal{L}_B = \frac{1}{JK} \sum_{j=1}^{J} \|\widehat{z}_\phi(e_j) - z_j^A\|_2^2 + \mathcal{L}_{\mathrm{cal}}(\widehat{z}_\phi) \\ + \frac{\lambda_{\mathrm{phys}}^B}{2} (\mathcal{L}_{\mathrm{phys}}^{\mathrm{obs}} + \mathcal{L}_{\mathrm{phys}}^{\mathrm{col}}) + \lambda_{\mathrm{sm}} \mathcal{L}_{\mathrm{sm}}. \end{aligned} \tag{15}$$

The code loss averages over environments and code dimensions. The physical term averages penalties at observed and collocation environments; calibration consistency is evaluated using predicted codes through the frozen modes.

A.3 Physical realizations and environment regularization

The operator interface combines the physical matrix with fixed source injection and sensor extraction:

$$G_{\mathrm{phys}}(e,\omega) = CL_{\mathrm{phys}}(e,\omega)^{-1}B. \tag{16}$$

For advection–diffusion, L_{phys} combines the finite-difference spatial matrix with the backward-Euler temporal symbol. The learned real correction is shared across frequencies and restricted to the supplied stencil. At zero temporal frequency, the physical loss penalizes nonpositive diagonal entries, positive off-diagonal entries, negative row sums, and an imaginary component in the corrected spatial operator. These are soft penalties on discrete operator properties; no separate neural solution field or pointwise PDE residual is introduced.

When the prior supplies only a transfer, we use a diagonal state indexed by source–receiver channels. The matrices B and C inject each source into its channels and collect them at the sensors. The diagonal physical matrix is constructed from regularized reciprocals of the supplied transfer, reproducing it up to the inverse ridge. Physical penalties discourage small corrected diagonal magnitudes and spectral curvature of the scaled correction. This realization supports effective propagation modeling without identifying a unique spatial differential operator.

The correction family is shared through the environment code, but its mode tensor follows the adapter's frequency structure. The advection–diffusion adapter ties the real correction modes across frequencies through one spatial realization, whereas the transfer-only OGW2 adapter stores frequency-indexed mode entries. Thus K is the shared code dimension; it does not imply that all frequency blocks have tied correction parameters. The corresponding tensor shapes and adapter settings are recorded in each run manifest.

In the perturbation results, G_{phys} denotes the transfer induced by this discrete realization. A transfer-only input prior G_{prior} can differ from it through reciprocal regularization. Comparing a corrected transfer with that input prior therefore requires the additional discrepancy $\|G_{\mathrm{phys}} - G_{\mathrm{prior}}\|_2$. The propositions' invertibility assumptions still apply: reciprocal regularization alone does not guarantee them at zero transfer entries. The diagnostic records distinguish the input prior and its realized transfer.

For Stage B, collocation covers a prescribed margin around calibration: evenly spaced centers for scalar environments and a scrambled Sobol design for vector environments. At centers $\{e_i\}_{i=1}^{I}$, let v_d be coordinate vector d and $h_d = 0.02 s_{e,d}$. The implemented smoothness loss is

$$\mathcal{L}_{\mathrm{sm}} = \frac{1}{Id_eK} \sum_{i=1}^{I} \sum_{d=1}^{d_e} \left\| \frac{\widehat{z}_\phi(e_i + h_d v_d) - 2\widehat{z}_\phi(e_i) + \widehat{z}_\phi(e_i - h_d v_d)}{\max(h_d^2, 10^{-8})} \right\|_2^2. \tag{17}$$

All centers contribute, including the endpoints; displaced evaluations can lie just outside the collocation interval. If collocation is disabled, calibration environments serve as centers.

B Conditional Error Analysis and Proofs

The results below concern the supplied discrete realization and frozen learned parameters. They are deterministic, conditional bounds built on perturbation and convex-recovery arguments (Higham, 2002; Lounici et al., 2011; Parikh & Boyd, 2014). They characterize the structure and sources of error, but they do not guarantee training convergence or identify a unique continuous PDE operator; finite-sample generalization is treated separately by the Rademacher bound in Section B.2. The reported relative Frobenius transfer errors are empirical quantities and are not used to estimate the spectral constants or restricted gains that appear below.

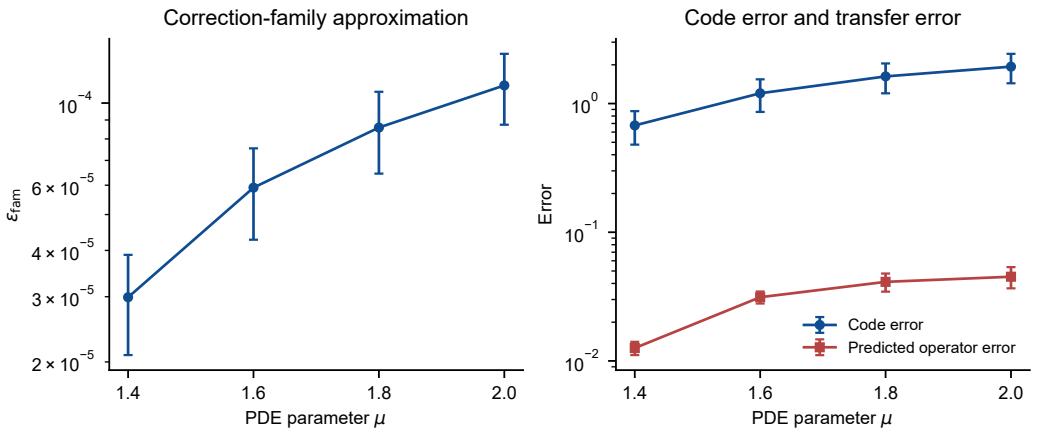


Figure 3: Transport bound-term decomposition. Left: correction-family approximation error ϵ_{fam} remains near 10^{-4} across OOD parameters. Right: environment-code error and predicted transfer error grow together with the PDE parameter. Points and error bars show mean $\pm$ sample SD across three PASL-U runs.

B.1 Operator stability and error decomposition

The entrywise budget in Eq. 3 controls correction magnitude. Stability of the inverse additionally requires a condition on the physical operator:

Proposition 1 (Bounded perturbation and conditional resolvent stability)**.** *For the parameterization in Eq. 10, let $\beta = \rho\|M \odot S\|_F$, $p = \|L_{\text{phys}}^{-1}\|_2$, and $\theta = p\beta$. If L_{phys} is invertible and $\theta < 1$, then every admissible code yields an invertible $\widehat{L}$ with*

$$\|\widehat{L}^{-1}\|_2 \le \frac{p}{1-\theta}, \qquad \|\widehat{G} - G_{\text{phys}}\|_2 \le \|C\|_2\|B\|_2\frac{p^2\beta}{1-\theta}. \tag{18}$$

For diagonal realizations, the tighter computable envelope $\beta_{\text{diag}} = \rho \max_i M_{ii}S_{ii}$ may replace β. Neither envelope is an empirical estimate of the learned inverse conditioning.

Proof of Proposition 1. For each matrix entry, $M \odot S$ is nonnegative and $D \ge |U_0| + \sum_k a_k|U_k|$. Thus, for any $z \in \mathcal{Z}$,

$$|\Delta L(z)| \le \rho M \odot S \odot \left(|U_0| + \sum_{k=1}^{K} a_k|U_k|\right) \oslash D \le \rho M \odot S. \tag{19}$$

Consequently $\|\Delta L(z)\|_2 \le \|\Delta L(z)\|_F \le \beta$. Factor the corrected operator as $\widehat{L} = L_{\text{phys}}(I + L_{\text{phys}}^{-1}\Delta L)$. When $\|L_{\text{phys}}^{-1}\Delta L\|_2 \le \theta < 1$, the Neumann series gives

$$\widehat{L}^{-1} = \sum_{m=0}^{\infty}(-L_{\text{phys}}^{-1}\Delta L)^m L_{\text{phys}}^{-1}, \qquad \|\widehat{L}^{-1}\|_2 \le p\sum_{m=0}^{\infty}\theta^m = \frac{p}{1-\theta}.$$

The resolvent identity $\widehat{L}^{-1} - L_{\text{phys}}^{-1} = -\widehat{L}^{-1}\Delta L L_{\text{phys}}^{-1}$ then yields

$$\|\widehat{G} - G_{\text{phys}}\|_2 \le \|C\|_2\|\widehat{L}^{-1}\|_2\|\Delta L\|_2\|L_{\text{phys}}^{-1}\|_2\|B\|_2,$$

which proves Eq. 18. For diagonal realizations, the spectral norm of ΔL equals its largest diagonal magnitude, so β can be replaced by $\rho \max_i M_{ii}S_{ii}$. A uniform bound over an environment and frequency set additionally requires a uniform gap from $\theta = 1$ and a uniform bound on p.

Proof of the error decomposition (Eq. 7). Suppress e, let $L_z = \widehat{L}(e, z)$, and write $d_\Delta = (\sum_k \|\Delta L_k\|_2^2)^{1/2}$. For $z, z' \in \mathcal{Z}$,

$$\widehat{G}(e,z) - \widehat{G}(e,z') = C L_z^{-1}(L_{z'} - L_z) L_{z'}^{-1} B,$$
$$\|L_{z'} - L_z\|_2 \le \sum_{k=1}^{K} |z'_k - z_k| \, \|\Delta L_k\|_2 \le d_\Delta \|z' - z\|_2. \tag{20}$$

The inverse bound at both codes implies $\|\widehat{G}(e,z) - \widehat{G}(e,z')\|_2 \le c_e \|z - z'\|_2$. Apply this inequality with $z = \widehat{z}_\phi(e)$ and $z' = z^\dagger(e)$, then use the triangle inequality with $G^\star(e)$ to obtain Eq. 7. A best-approximating code exists for a single transfer target because the box is compact and the transfer is continuous under the stated invertibility assumption. When a code is shared across frequencies, one may instead choose it to minimize an aggregate discrepancy; the pointwise bound still holds with its corresponding residual at each frequency. The argument requires no latent uniqueness and no true differential operator with the same internal realization.

The differential form also exposes the role of applying corrections before inversion:

$$\frac{\partial \widehat{G}(e,z)}{\partial z_k} = -C \widehat{L}(e,z)^{-1} \Delta L_k \widehat{L}(e,z)^{-1} B. \tag{21}$$

Thus, mode size and resolvent conditioning jointly determine the sensitivity to each coordinate. Proposition 1 supplies the sufficient choice $\kappa_e = p/(1-\theta)$ when its condition holds at the frequency under consideration.

B.2 Finite-sample environment-code generalization

Proposition 2 (Finite-sample environment-code generalization)**.** *Let $\mathcal{E} \subset \mathbb{R}^{d_e}$ be compact, and let calibration environments $e_1, \dots, e_J$ be independent and identically distributed on $\mathcal{E}$. Suppose the environment-code predictor $\widehat{z}_\phi : \mathcal{E} \to \mathcal{Z}$ is L_ϕ-Lipschitz, the comparison code map $z^\dagger$ is $L_\dagger$-Lipschitz, and the code loss $\ell(e) = \|\widehat{z}_\phi(e) - z^\dagger(e)\|_2$ is bounded by M_ℓ. Let $\mathcal{R}_J(\Phi)$ denote the empirical Rademacher complexity of the predictor class. Then, with probability at least $1 - \delta$,*

$$\sup_{e \in \mathcal{E}} \|\widehat{z}_\phi(e) - z^\dagger(e)\|_2 \le \frac{1}{J} \sum_{j=1}^{J} \|\widehat{z}_\phi(e_j) - z^\dagger(e_j)\|_2 + 2\mathcal{R}_J(\Phi) + M_\ell \sqrt{\frac{\log(1/\delta)}{2J}}. \tag{22}$$

The bound separates three contributions: the empirical calibration-code error, a capacity term controlled by the predictor class, and the usual confidence term, which decays as $O(J^{-1/2})$. For a query e outside the empirical support, Eq. 22 is evaluated at a reference calibration environment e_b, and the Lipschitz relation in Eq. 8 adds the distance term $(L_\phi + L_\dagger)\|e - e_b\|_2$. This yields a high-probability bound in which code error grows with both finite-sample estimation error and environmental distance.

Proof. Let $\mathcal{F} = \{e \mapsto \|\widehat{z}_\phi(e) - z^\dagger(e)\|_2 : \phi \in \Phi\}$ be the loss class and $\widehat{R}_J(\mathcal{F})$ its empirical Rademacher complexity. Define the empirical error

$$\widehat{L}_J(\phi) = \frac{1}{J} \sum_{j=1}^{J} \|\widehat{z}_\phi(e_j) - z^\dagger(e_j)\|_2.$$

Changing one calibration environment changes $\widehat{L}_J(\phi)$ by at most M_ℓ / J, so McDiarmid's inequality gives, with probability at least $1 - \delta$,

$$\sup_{\phi \in \Phi} \Big(L(\phi) - \widehat{L}_J(\phi) \Big) \le \mathbb{E}\left[\sup_{\phi \in \Phi} \Big(L(\phi) - \widehat{L}_J(\phi) \Big) \right] + M_\ell \sqrt{\frac{\log(1/\delta)}{2J}},$$

where $L(\phi) = \sup_{e \in \mathcal{E}} \|\widehat{z}_\phi(e) - z^\dagger(e)\|_2$. By the standard symmetrization inequality, the expectation is bounded by $2\widehat{R}_J(\mathcal{F})$. Since the Euclidean norm is 1-Lipschitz and $z^\dagger$ is a fixed comparison map, Talagrand's contraction principle gives $\widehat{R}_J(\mathcal{F}) \le \mathcal{R}_J(\Phi)$. Combining these steps yields Eq. 22. The outside-support statement then follows by applying Eq. 8 from e_b to e. □

Across the three transport runs, the empirical Lipschitz constants are $L_\phi \in [3.58, 4.30]$ and $L_\dagger \in [3.96, 4.48]$, so the distance term in Eq. 9 is quantitatively active in the tested OOD range. The Rademacher term $\mathcal{R}_J(\Phi)$ is not Monte-Carlo estimated here and is left as the standard capacity term for the Lipschitz MLP class.

The bound remains conditional on the invertibility and conditioning assumptions used to control c_e in Eq. 9.

B.3 Source recovery error

Let $\widehat{A}, A^\star : \mathbb{R}^{QT} \to \mathbb{R}^{PT}$ be the vectorized time-domain actions of the predicted and true operators at deployment. The background-subtracted observation is $r_\star = A^\star s^\star + \eta = \widehat{A}s^\star + \nu$, where

$$\eta = \varepsilon + b - \widehat{b}, \qquad \nu = (A^\star - \widehat{A})s^\star + \eta, \qquad \|\nu\|_2 \le \|A^\star - \widehat{A}\|_2 \|s^\star\|_2 + \|\eta\|_2. \tag{23}$$

For the real rFFT/irFFT action, $\|\widehat{A} - A^\star\|_2 \le \max_f \|\widehat{G}_f - G_f^\star\|_2$. At DC and Nyquist frequencies, the effective real blocks are the real parts of the supplied matrices. Parseval's identity gives this inequality, connecting the frequency-wise decomposition in Eq. 7 to the source-dependent effective noise.

Regularized group-sparse recovery.

Proposition 3 (Conditional group-sparse recovery)**.** *Suppose $s^\star$ has s_0 active groups indexed by $\mathcal{I} \subseteq \{1, \dots, Q\}$, and $\widehat{s}$ minimizes Eq. 6. Let $\mathcal{I}^c$ be the complement of $\mathcal{I}$, and let $h_\mathcal{I}$ and $h_{\mathcal{I}^c}$ denote the corresponding group restrictions. Write $\|h\|_{2,1} = \sum_q \|h_q\|_2$ and $\|v\|_{2,\infty} = \max_q \|v_q\|_2$. Assume $\lambda_s > 0$, $\lambda_s \ge 2\|\widehat{A}^* \nu\|_{2,\infty}$, and*

$$\|\widehat{A}h\|_2 \ge \alpha \|h\|_2 \quad \textit{whenever} \quad \|h_{\mathcal{I}^c}\|_{2,1} \le 3\|h_\mathcal{I}\|_{2,1}, \qquad \alpha > 0.$$

Then

$$\|\widehat{s} - s^\star\|_2 \le \frac{3\lambda_s \sqrt{s_0}}{\alpha^2}. \tag{24}$$

Proof of Proposition 3. Let $h = \widehat{s} - s^\star$, abbreviate λ_s as λ, and write $\widehat{A}$ as A. Optimality of $\widehat{s}$ relative to the feasible comparison $s^\star$ implies

$$\tfrac{1}{2}\|Ah - \nu\|_2^2 + \lambda\|s^\star + h\|_{2,1} \le \tfrac{1}{2}\|\nu\|_2^2 + \lambda\|s^\star\|_{2,1}.$$

After expanding the square and using group-norm duality,

$$\begin{aligned} \tfrac{1}{2}\|Ah\|_2^2 &\le \mathrm{Re}\langle Ah, \nu\rangle + \lambda(\|s^\star\|_{2,1} - \|s^\star + h\|_{2,1}) \\ &\le \tfrac{\lambda}{2}\|h\|_{2,1} + \lambda(\|h_\mathcal{I}\|_{2,1} - \|h_{\mathcal{I}^c}\|_{2,1}) \\ &= \tfrac{3\lambda}{2}\|h_\mathcal{I}\|_{2,1} - \tfrac{\lambda}{2}\|h_{\mathcal{I}^c}\|_{2,1}. \end{aligned} \tag{25}$$

Nonnegativity of the left side and $\lambda > 0$ give the cone condition. The restricted gain assumption and Cauchy–Schwarz across the s_0 active groups then imply

$$\alpha^2\|h\|_2^2 \le \|Ah\|_2^2 \le 3\lambda\|h_\mathcal{I}\|_{2,1} \le 3\lambda\sqrt{s_0}\|h\|_2.$$

For $h \neq 0$, division by $\alpha^2\|h\|_2$ proves Eq. 24; the zero case is immediate.

For a group-error bound $E_s = 3\lambda_s\sqrt{s_0}/\alpha^2$, every individual group error is at most E_s. If the weakest active group has norm greater than $2E_s$, its estimated norm exceeds that of every inactive group, so selecting the s_0 largest groups recovers the support. This sufficient separation condition illustrates why support recovery depends on source strength as well as operator mismatch and inverse conditioning.

Finite-step least squares. With $\lambda_s = 0$, proximal gradient reduces to gradient descent. The following result describes its actual finite-step output without a full-rank assumption.

Proposition 4 (Zero-initialized finite-step recovery)**.** *For a fixed nonzero $\widehat{A}$ and $0 < \tau \le \|\widehat{A}\|_2^{-2}$, let $s^{(0)} = 0$ and*

$$s^{(t+1)} = s^{(t)} - \tau \widehat{A}^*(\widehat{A}s^{(t)} - r_\star).$$

Define $R_\tau = I - \tau \widehat{A}^*\widehat{A}$ and $F_t = \tau \sum_{j=0}^{t-1} R_\tau^j \widehat{A}^*$. For any integer $t \ge 1$,

$$s^{(t)} - s^\star = -R_\tau^t s^\star + F_t(A^\star - \widehat{A})s^\star + F_t\eta. \tag{26}$$

If $\sigma_i > 0$ are the nonzero singular values of $\widehat{A}$, set

$$g_t = \max_{\sigma_i > 0} \frac{1 - (1 - \tau\sigma_i^2)^t}{\sigma_i}.$$

Then

$$\|s^{(t)} - s^\star\|_2 \le \underbrace{g_t\|A^\star - \widehat{A}\|_2\|s^\star\|_2}_{\text{operator mismatch}} + \underbrace{g_t\|\eta\|_2}_{\text{observation error}} + \underbrace{\|R_\tau^t s^\star\|_2}_{\text{iteration and null-space bias}}. \tag{27}$$

Proof. Substituting $r_\star = \widehat{A}s^\star + \nu$ into the iteration gives $s^{(t+1)} - s^\star = R_\tau(s^{(t)} - s^\star) + \tau\widehat{A}^*\nu$. Unrolling from zero proves Eq. 26. A singular value decomposition of $\widehat{A}$ shows that the nonzero singular values of F_t are $[1 - (1 - \tau\sigma_i^2)^t]/\sigma_i$, so $\|F_t\|_2 = g_t$. The triangle inequality proves Eq. 27.

Let P_0 be the orthogonal projector onto $\ker \widehat{A}$, let $\sigma_{\min}^+$ be its smallest positive singular value, and set $q_\tau = 1 - \tau(\sigma_{\min}^+)^2 \in [0, 1)$. Orthogonality of the null space and its complement gives

$$\|R_\tau^t s^\star\|_2^2 \le \|P_0 s^\star\|_2^2 + q_\tau^{2t}\|(I - P_0)s^\star\|_2^2. \tag{28}$$

For full column rank, $P_0 = 0$ and this term decays geometrically. With rank deficiency, the unobservable source component remains. Meanwhile, g_t is nondecreasing and bounded by $1/\sigma_{\min}^+$, so additional iterations can amplify mismatch and observation error even as the recoverable bias decreases. If $\widehat{A} = 0$, every iterate is zero and the recovery error is $\|s^\star\|_2$.

At 30 dB, transport uses $\lambda_s = 0$ and $t = 120$. Its Frobenius-based step (Appendix C.2) satisfies $\tau\|\widehat{A}\|_2^2 \le 0.9$, so the bound applies to these source-linear solvers in exact arithmetic. It explains why operator-error and finite-step recovery rankings can differ, without estimating the separate terms from the reported metrics. Lunz-FAC's learned direction need not equal an adjoint and is outside this result.

Finite regularized iterates. Proposition 3 concerns an exact minimizer. A finite iterate additionally incurs $\|s^{(t)} - \widehat{s}\|_2$. Under its assumptions and the additional condition of full column rank with smallest singular value $\sigma_{\min} > 0$, proximal-gradient descent with $0 < \tau \le \|\widehat{A}\|_2^{-2}$ satisfies

$$\|s^{(t)} - s^\star\|_2 \le \frac{3\lambda_s\sqrt{s_0}}{\alpha^2} + \frac{\|s^{(0)} - \widehat{s}\|_2}{\sigma_{\min}\sqrt{\tau t}}.$$

Indeed, the objective gap is at most $\|s^{(0)} - \widehat{s}\|_2^2/(2\tau t)$ (Parikh & Boyd, 2014), while strong convexity lower-bounds it by $\sigma_{\min}^2\|s^{(t)} - \widehat{s}\|_2^2/2$. The cone condition alone does not give this distance bound for arbitrary iterates. Neither full column rank nor dual-noise domination is verified by the guided-wave experiments, so their positive-regularization results are not certified by this bound.

Profiled background as a weighted problem. For a frozen intact template $\widehat{b}$, the synthetic OGW2 core solver profiles a scalar background amplitude with ridge $\gamma_b = 0.02$. For vectorized observations, eliminating that amplitude from

$$\min_{s,\xi} \ \tfrac{1}{2}\|\widehat{A}s + \xi\widehat{b} - y\|_2^2 + \tfrac{\gamma_b}{2}\xi^2 + \lambda_s\|s\|_{2,1}$$

gives the residual weight $Q_b = I - \widehat{b}\widehat{b}^*/(\|\widehat{b}\|_2^2 + \gamma_b)$. The implemented gradient is $\widehat{A}^*Q_b(\widehat{A}s - y)$. Let $W_b = Q_b^{1/2}$ be the positive-definite matrix square root, $\bar{A} = W_b\widehat{A}$, and $\bar{y} = W_b y$. The profiled objective is $\frac{1}{2}\|\bar{A}s - \bar{y}\|_2^2 + \lambda_s\|s\|_{2,1}$. Its effective noise is

$$\bar{\nu} = W_b\big[(A^\star - \widehat{A})s^\star + b + \varepsilon\big].$$

Proposition 3 applies with $(\widehat{A}, \nu, \alpha)$ replaced by $(\bar{A}, \bar{\nu}, \bar{\alpha})$, provided that the restricted-gain assumption holds for $\bar{A}$ and

$$\lambda_s \geq 2\big\|\widehat{A}^* Q_b[(A^\star - \widehat{A})s^\star + b + \varepsilon]\big\|_{2,\infty}.$$

Thus the assumptions concern the weighted inverse problem, including the selected background template. For any reference amplitude $\xi^\star$, set $d_b = b - \xi^\star \widehat{b}$. Then

$$\|\bar{\nu}\|_2 \leq \|W_b(A^\star - \widehat{A})\|_2 \|s^\star\|_2 + \|W_b(\varepsilon + d_b)\|_2 + |\xi^\star| \sqrt{\frac{\gamma_b}{\|\widehat{b}\|_2^2 + \gamma_b}}\, \|\widehat{b}\|_2.$$

The final term is ridge-induced background residual, which can remain even with the true template. For positive ridge, Q_b is a positive-definite contraction and generally is not idempotent. The eliminated quadratic is

$$\tfrac{1}{2}(\widehat{A}s - y)^* Q_b(\widehat{A}s - y).$$

If regularization is zero, Proposition 4 applies with action $\bar{A}$, true action $W_b A^\star$, and observation error $W_b(b + \varepsilon)$. For positive regularization, the finite-iterate extension above additionally requires full column rank of $\bar{A}$. Since $\|\bar{A}\|_2 \leq \|\widehat{A}\|_2$, the unweighted step bound is also valid for the weighted problem. Without a template, $Q_b = I$.

C Experimental Details

C.1 Datasets and evaluation protocol

Advection–diffusion. The controlled transport system follows

$$\partial_t u - \nabla \cdot (\kappa(x, \mu) \nabla u) + v(\mu) \cdot \nabla u + c(x, \mu) u = s(x, t). \tag{29}$$

We use zero Dirichlet boundary conditions on the outer boundary, which is not observed. The discretization contains a 21×21 grid, 64 internal sensors, 25 candidate source locations, and 32 time samples with step size 0.08. The data-generating system and the physical prior share this PDE class but have different smooth parameter-dependent diffusion, velocity, and reaction laws. The prior realization combines a finite-difference spatial matrix with the backward-Euler temporal symbol, as described in Appendix A.3.

The fitting parameters are $\{0.6, 0.65, 0.75, 0.8, 0.85, 1.0, 1.05, 1.15, 1.2\}$, with 128 probes per parameter. We reserve $\mu = 1.1$ for validation, $\{0.7, 0.9\}$ for interpolation tests, and $\{1.4, 1.6, 1.8, 2.0\}$ for extrapolation. Test sources and true test operators are used only to compute evaluation metrics.

Dual-parameter transport. The two-parameter system uses the same finite-difference geometry, source/sensor layout, and time discretization as the single-parameter transport benchmark. Calibration uses the 16 grid environments $c, d \in \{0.8, 0.95, 1.1, 1.2\}$, and evaluation uses the eight extrapolation environments $[1.6, 1.0]$, $[2.0, 1.0]$, $[1.0, 1.6]$, $[1.0, 2.0]$, $[1.4, 1.4]$, $[1.6, 1.6]$, $[1.8, 1.8]$, and $[2.0, 2.0]$. Each calibration environment contains 16 probes at 40 dB. Each test environment contains 80 single-source and 80 two-source examples at each noise level (20, 30, and 40 dB); main comparisons use 30 dB. Training uses 2000 epochs with early-stopping patience 100. PASL-U uses rank 6, hidden width 16, learning rate 5×10^{-3}, correction ratio 0.5, latent bound 2.0, action loss weight 0.1, prior loss weight 10^{-3}, and physics/smoothness weights $0.01/0.05/10^{-3}$. Data-NGO, Neumann-NGO, and Lunz-FAC use learning rate 10^{-3} and hidden width 16; PINO-style uses rank 8, hidden width 32, and prior weight 0.25; DeepONet and FNO use hidden width 32 and latent dimension 16. Recovery uses group-sparse reconstruction with regularization 0.05 and 150 iterations.

Synthetic guided waves. Calibration temperatures range from 20 to $45°$C in increments of $2.5°$C. Each temperature contains 20 scattering probes and five direct-action records, split 80/20 for fitting and internal validation. The system has 66 measurement paths, four candidate sources, and 128 time samples at 128 kHz. Calibration SNR is 40 dB. At each test temperature and SNR, a run contains 200 single-source and 120 two-source examples. The test temperatures are 50, 55, and $60°$C. SNR is defined relative to the scattering signal before adding the intact response. The same frozen PASL-U models are evaluated at 20, 30, and 40 dB; the main comparisons use 30 dB.

Measured guided waves. The measured dataset contains 66 pitch–catch paths at a 40 kHz carrier frequency. Responses are cropped to 0.3–1.3 ms, reference pitch signals to 0–1 ms, and both are mean-detrended and resampled to 256 points. The four damage-location classes D04, D12, D16, and D24 correspond to candidate coordinates $(65, 400)$, $(195, 330)$, $(335, 260)$, and $(450, 190)$ mm in the dataset coordinate system.

Operator calibration uses 194 intact records with $20 \leq T < 45°\mathrm{C}$. Damage records below $45°\mathrm{C}$ are split by class and acquisition date into 318 classifier-training and 82 validation examples, with no record or batch overlap between these sets. Damage labels train the classifier but do not update propagation. Only the below-$45°\mathrm{C}$ validation partition selects classifier regularization. The three higher-temperature bands contain 80, 80, and 84 records, as reported in Table 4. These higher-temperature records were available during historical development; the reported experiment is a fixed-protocol re-evaluation rather than an independent blind acquisition test.

Table 4: Measured damage-class counts under split seed 20260903. The 194 intact calibration records are additional to these damage records.

Partition	D04	D12	D16	D24	Total
Training, $T < 45°\mathrm{C}$	82	80	77	79	318
Validation, $T < 45°\mathrm{C}$	18	20	23	21	82
Evaluation, $45 \leq T < 50°\mathrm{C}$	20	20	20	20	80
Evaluation, $50 \leq T < 55°\mathrm{C}$	20	20	20	20	80
Evaluation, $55 \leq T \leq 60°\mathrm{C}$	21	21	21	21	84

Repeated runs. We use seeds 20260903, 20260904, 20260905, 20260906, and 20260907. Synthetic guided-wave runs vary the simulated probes, sources, noise, fitting initialization, and within-temperature split; methods receive the same observations within each run. Advection–diffusion uses one fixed dataset and varies fitting initialization. Measured OGW2 uses the fixed split seed 20260903 and varies model training. The reported variation therefore has a different scope in each benchmark and does not measure variability across independent measurement campaigns.

Metrics. For an environment e, operator and source errors are

$$\mathrm{RelErr}(G_e) = \frac{\|\widehat{G}_e - G_e^\star\|_F}{\|G_e^\star\|_F}, \qquad \mathrm{NMSE}(\widehat{s}, s) = \frac{\|\widehat{s} - s\|_2^2}{\|s\|_2^2}. \tag{30}$$

Transfer blocks are stacked over the frequencies or time-domain channels used by each benchmark. Source tensors are vectorized for NMSE, so the numerator is the total squared waveform error. Top-1 ranks candidate groups by their recovered waveform norms. Top-2 recall is the fraction of the two active groups appearing among the two largest recovered groups. For the four-candidate guided-wave task, random ranking gives Top-1 accuracy 25% and Top-2 recall 50%; a zero prediction has NMSE 1 for any nonzero source.

We first average over examples within each environment and run. Interval summaries then average the environment-level values within a run, followed by the mean and sample SD across the five runs. Temperature-resolved entries omit the environment-averaging step. We do not pool test examples across seeds to estimate the SD. Measured recognition is evaluated by accuracy and Macro-F1. Accuracy equals balanced accuracy in each evaluation band because the four classes have equal counts. In method-comparison tables, bold marks the best non-Oracle mean for each condition, including ties at the displayed precision.

C.2 Implementation details

Physical baseline. Physical-only evaluates the supplied environment-dependent prior without a learned correction. Recovery comparisons across all three benchmarks include Physical-only, Data-NGO, Neumann NGO, Lunz-FAC, PINO-style, DeepONet, and FNO.

PINO-style. PINO-style predicts reduced transfer coefficients from calibration probes via a physics-regularized Fourier network. Following the use of physical information to constrain operator learning (Li et al., 2024), its objective adds a coefficient-space penalty toward the supplied

physical transfer, with weight 0.25. It uses AdamW with learning rate 0.002, weight decay 10^{-6}, and 2000 updates per fit; the lowest calibration training objective selects the checkpoint. Transport uses all calibration probes and the same validation-only recovery-regularization search. Synthetic waves use the same 80/20 calibration split and fit scattering and intact responses separately. Measured waves use all 194 intact calibration records, averaging direct-transfer targets at repeated temperatures before fitting. Its predicted direct response supplies background subtraction and the shared candidate-geometry lift below. Neither this method nor the other learned baselines uses test observations or true test operators for fitting.

Data-NGO. Data-NGO learns a source-linear Green's map from calibration transfer snapshots (Melchers et al., 2026). It uses a bias-free residual system network with three hidden layers, width 32 for guided waves and width 8 for transport.

Neumann NGO. Neumann NGO uses the physical ansatz of Melchers et al. (2026), Eq. (4.16), truncated at order one:

$$R(e) = -(L_{\mathrm{phys}}(e) - L_0)L_0^{-1}, \qquad \widehat{G}(e) = CL_0^{-1}[I + R(e) + \mathrm{NN}(R(e))]B.$$

The reference is $30°$C or $\mu = 1$. The bias-free residual network has three width-32 hidden layers, real/imaginary channels, and weights shared across frequencies. Transport uses a dense realization and guided waves a diagonal channel realization, with PASL-U's L_{phys}, B, C. Calibration-prior features determine normalization. Training uses all calibration snapshots and responses, Adam at 10^{-3}, and 2000 updates. Held-out wave probes select checkpoints; transport and measured data use calibration training loss. The finite ansatz is evaluated without test-dependent clipping.

Lunz-FAC. The forward–adjoint correction baseline (Lunz et al., 2021) learns

$$A_\theta(s, e) = F_\theta(A_{\mathrm{phys}}(e)s, e), \qquad D_\phi(r, e) = G_\phi(A_{\mathrm{phys}}(e)^* r, e).$$

Both corrections are zero-preserving residual MLPs with three width-32 hidden layers and environmental, frequency, and real/imaginary inputs. The forward correction mixes sensor channels; the learned direction applies a shared coordinatewise source correction. Inputs are normalized by calibration-only RMS values. Identified calibration operators provide forward and adjoint targets on probes and detached reconstruction trajectories whose length increases from one to ten iterates. The implementation uses signed waveforms, zero initialization, and group sparsity; its coordinatewise source correction has a different expressivity from the original imaging architecture.

Each fit uses 2000 joint Adam updates, learning rate 10^{-3}, and batch size 16. Every 100 updates, a ten-iterate calibration loss selects the checkpoint using held-out probes for synthetic waves and calibration-training records for the other tasks. Recovery takes a proximal step along the learned forward–adjoint direction, with step fraction 0.9 relative to the physical transfer's squared Frobenius bound. This direction need not be an exact gradient, so Lunz-FAC is evaluated through source recovery rather than a single source-linear transfer matrix.

Candidate geometry for measured data. PASL-U, Data-NGO, Neumann NGO, and PINO-style apply learned path factors to the same physical candidate-source transfer. For Data-NGO and the reduced-coefficient method, the predicted direct response divided by the physical direct response multiplies each candidate column, using the same zero guard as Neumann NGO. Physical-only retains the physical transfer, and Lunz-FAC combines its learned corrections with the same geometry. Intact records constrain direct propagation paths; transferring these corrections to candidate-source columns is a modeling assumption.

PASL-U training. The environment map is an MLP with two width-32 hidden layers. We use correction budget $\rho = 0.5$, code bounds $a_k = 2$, physical-loss weights $\lambda^A_{\mathrm{phys}} = 0.01$ and $\lambda^B_{\mathrm{phys}} = 0.05$, and smoothness weight $\lambda_{\mathrm{sm}} = 0.001$. Direct-action records, including measured calibration, use the closed-form channel inverse with ridge $\eta = 10^{-4}$ described in Appendix A.2. Table 5 reports the CUDA optimization budgets. Each Stage-A or Stage-B iteration performs one optimizer update; the classifier uses full-data epochs. Neumann NGO and Lunz-FAC use 2000 updates per fit, so total optimization work is method-dependent.

Table 5: Training settings for the core experiments. Stage columns show completed optimizer updates; collocation intervals specify centers before smoothness offsets.

Fit	K	(e_0, s_e)	Stage A	Stage B	Collocation range
Synthetic OGW2 scattering	6	$(30, 10)$	2000	2000	$[13.75, 51.25]$
Synthetic OGW2 direct action	4	$(30, 10)$	2000	2000	$[13.75, 51.25]$
Advection–diffusion transport	6	$(1, 0.4)$	2000	2000	$[0.45, 1.35]$
Measured OGW2 transfer	8	$(30, 10)$	2000	2000	$[14.51, 50.79]$

Stage B uses 16 unlabeled collocation centers. The prior and predicted operator are evaluated there without supplying measured targets or true test transfers. The upper collocation limit lies below every transport OOD parameter. For synthetic guided waves it extends beyond calibration into the first test-temperature region, and for measured data it reaches the near and lower-middle evaluation bands. Smoothness evaluations use coordinate offsets of $0.02 s_e$ around these centers, as defined in Eq. 17.

Source inversion. The source-linear methods initialize the source at zero and use a fixed number of proximal-gradient steps with the frozen predicted operator and its exact adjoint. For frequency blocks $\widehat{G}_f$, the step is

$$\tau = \frac{0.9}{\max\{\max_f \|\widehat{G}_f\|_F^2, 10^{-8}\}}.$$

This denominator is based on a Frobenius upper bound on the squared operator norm. We use fixed iteration counts rather than a tolerance-based stopping rule. Transport uses 120 iterations and selects λ_s separately for each method and SNR at $\mu = 1.1$ from $\{0, 0.001, 0.003, 0.01, 0.03, 0.1\}$. The reported source-linear methods select zero at 30 dB. Synthetic and measured guided waves use 100 and 250 iterations, respectively, with fixed $\lambda_s = 0.1$. Lunz-FAC uses the same iteration counts and regularization protocol with its learned direction.

Background estimation. For synthetic guided waves, scattering probes calibrate source transfer and intact direct-action records calibrate background prediction. Recovery profiles a scalar background amplitude with ridge $\gamma_b = 0.02$, using the objective and gradient in Appendix B.3. Measured recovery instead subtracts the predicted intact response, obtained by passing each record's pitch-derived reference input through the corrected direct-response model, without amplitude profiling. The synthetic-wave true-transfer reference retains the estimated PASL-U background while replacing only the source transfer.

The true-operator references use the same finite-iteration recovery rule as the source-linear methods and are not theoretical lower bounds on NMSE. The transport results with $\lambda_s = 0$ correspond to Proposition 4; positive-regularization recovery is analyzed separately in Appendix B.3.

Measured feature construction. For each record, the frozen inverse returns a real source tensor $\widehat{s} \in \mathbb{R}^{4\times 256}$. We flatten its signed waveform values directly into 1024 features. Each feature is centered and scaled using its mean and standard deviation over the 318 classifier-training damage records. Standard deviations are floored at 10^{-6}. The same statistics transform validation and evaluation records.

Linear classifier. A linear softmax head maps 1024 features to the four damage classes defined in Appendix C.1. For each weight decay in $\{0, 10^{-5}, 10^{-4}, 10^{-3}, 10^{-2}\}$, AdamW trains for 2000 epochs with learning rate 3×10^{-3} and batch size 64. Cross-entropy uses training-class weights $N/(4N_c)$. Selection maximizes validation balanced accuracy, then Macro-F1, then prefers smaller weight decay. Each method fits separate classifier weights under this same protocol; propagation parameters remain frozen.

Checkpoint reuse. Advection–diffusion interpolation, OOD evaluation, and the core source-recovery results use the same frozen PASL-U checkpoints. Data-NGO also uses one frozen model per seed across interpolation and extrapolation.

C.3 COMPONENT ABLATION

We ablate the core PASL-U components on the dual-parameter advection–diffusion extrapolation task, using five seeds and the same recovery protocol as the main experiment. Table 6 reports OOD-averaged operator error and two-source NMSE, together with the two-source NMSE at the farthest environment $[2.0, 2.0]$.

Table 6: Core-module ablation on the dual-parameter transport task, five-seed mean $\pm$ sample SD. OOD columns average over the eight extrapolation environments.

Variant	OOD operator error	OOD two-source NMSE	$[2.0, 2.0]$ two-source NMSE
Full PASL-U	0.0510 ± 0.0059	0.2758 ± 0.0089	0.2778 ± 0.0077
No physics anchor	0.1597 ± 0.0062	0.3196 ± 0.0085	0.3835 ± 0.0105
No Stage B	0.1031 ± 0.0001	0.3122 ± 0.0075	0.3184 ± 0.0078
Joint single-stage	0.0529 ± 0.0049	0.2745 ± 0.0059	0.2766 ± 0.0077

The ablation results support the distinct roles of the physics anchor and environment-code prediction. Removing the physics anchor substantially increases both operator and recovery error, indicating that the physical prior provides an important reference for the learned correction. Removing Stage B also increases OOD error because a fixed reference code cannot represent environment-dependent changes in the correction. In contrast, jointly optimizing the correction and environment map achieves nearly the same accuracy as the explicit two-stage procedure. This indicates that the two-stage formulation primarily serves to separate the two error sources in Eq. 7 rather than being necessary for predictive accuracy.

C.4 STRICT OOD ADVECTION–DIFFUSION ADDITIONAL RESULTS

Table 7 gives interval-averaged transfer errors for advection–diffusion, and Table 8 gives parameter-resolved single-source OOD recovery.

Table 7: Advection–diffusion relative transfer-operator error under strict OOD extrapolation, five-seed mean $\pm$ sample SD (lower is better).

Method	$\mu = 1.4$	$\mu = 1.6$	$\mu = 1.8$	$\mu = 2.0$
Physical-only	0.0853 ± 0.0000	0.0927 ± 0.0000	0.0913 ± 0.0000	0.0871 ± 0.0000
PASL-U	$\mathbf{0.0119 \pm 0.0005}$	$\mathbf{0.0301 \pm 0.0014}$	$\mathbf{0.0502 \pm 0.0031}$	$\mathbf{0.0655 \pm 0.0054}$
Data-NGO	0.2076 ± 0.0710	0.4010 ± 0.1632	0.6725 ± 0.2708	1.0082 ± 0.3967
Neumann NGO	0.0419 ± 0.0083	0.0986 ± 0.0384	0.2185 ± 0.0798	0.4154 ± 0.1516
PINO-style	0.4670 ± 0.0969	0.5714 ± 0.0856	0.5325 ± 0.1128	0.9036 ± 0.0652

Table 8: Advection–diffusion single-source NMSE at 30 dB under strict OOD extrapolation, five-seed mean $\pm$ sample SD (lower is better).

Method	$\mu = 1.4$	$\mu = 1.6$	$\mu = 1.8$	$\mu = 2.0$
Physical-only	0.5587 ± 0.0000	0.5391 ± 0.0000	0.5281 ± 0.0000	0.5132 ± 0.0000
PASL-U	0.5238 ± 0.0005	0.4919 ± 0.0007	$\mathbf{0.4735 \pm 0.0015}$	$\mathbf{0.4534 \pm 0.0026}$
Data-NGO	0.5716 ± 0.0303	0.5907 ± 0.0556	0.6409 ± 0.0621	0.6885 ± 0.0633
Neumann NGO	$\mathbf{0.5206 \pm 0.0036}$	$\mathbf{0.4839 \pm 0.0158}$	0.4747 ± 0.0379	0.5224 ± 0.0799
Lunz-FAC	0.5600 ± 0.0007	0.5411 ± 0.0008	0.5305 ± 0.0014	0.5162 ± 0.0017
PINO-style	0.6859 ± 0.0255	0.6963 ± 0.0225	0.6751 ± 0.0352	0.7503 ± 0.0125
DeepONet	0.9131 ± 0.0024	0.9186 ± 0.0014	0.9365 ± 0.0016	0.9233 ± 0.0019
FNO	1.0095 ± 0.0051	1.0224 ± 0.0072	1.0390 ± 0.0126	1.0398 ± 0.0244

In single-source recovery, PASL-U also attains the lowest OOD-average NMSE (0.4884) among the non-Oracle methods, lower than Physical-only (0.5348) and Neumann NGO (0.5210), while FNO and DeepONet remain close to 1. The ordering observed in two-source recovery therefore also holds for single-source recovery.

C.5 OPERATOR-ERROR ANALYSIS

Table 9 reports the propagation-operator relative error by extrapolation environment for the methods that produce an explicit transfer operator.

Table 9: Propagation-operator relative error by extrapolation environment, five-seed mean (lower is better).

Environment	PASL-U	Physical-only	PINO-style	Data-NGO	Neumann-NGO
$[1.6, 1.0]$	$\mathbf{0.0403}$	0.1251	0.2445	0.4201	0.2568
$[2.0, 1.0]$	$\mathbf{0.0703}$	0.1224	0.3477	0.8517	0.6559
$[1.0, 1.6]$	$\mathbf{0.0326}$	0.0689	0.1572	0.3468	0.0446
$[1.0, 2.0]$	0.0613	0.0841	0.2232	0.5491	$\mathbf{0.0445}$
$[1.4, 1.4]$	$\mathbf{0.0230}$	0.1310	0.1752	0.3886	0.1322
$[1.6, 1.6]$	$\mathbf{0.0435}$	0.1459	0.4436	0.6607	0.2983
$[1.8, 1.8]$	$\mathbf{0.0607}$	0.1542	0.7510	1.0521	0.5123
$[2.0, 2.0]$	$\mathbf{0.0764}$	0.1613	0.5965	1.5277	0.7787

The operator-error results are consistent with the source-recovery results: PASL-U has the lowest error at every environment (0.023 to 0.076) and grows the slowest. Data-NGO and PINO-style degrade fastest as both parameters increase, while Neumann-NGO remains low only when d is extrapolated but rises sharply once c also increases. This confirms that the propagation-operator correction is the main source of PASL-U's stable extrapolation.

C.6 SYNTHETIC GUIDED-WAVE SOURCE RECOVERY

We consider temperature-dependent guided waves with 66 measurement paths and four candidate source locations. Observations combine source-induced scattering, an intact-structure response, and additive noise. We calibrate the framework at 20–45°C and recover unknown sources at 50, 55, and 60°C. Propagation prediction, background estimation, and inversion are evaluated together as components of the recovery pipeline.

Within calibration, held-out scattering probes give source NMSE 0.095 ± 0.003. These probes are used for checkpoint selection and provide a calibration-validation reference. From Table 11, at the nearest extrapolation temperature 50°C, PASL-U attains the lowest single-source NMSE (0.313) with 100% Top-1 and Top-2 recall, retaining both waveform and location information. As temperature moves farther from calibration, all methods degrade, but PASL-U degrades the most slowly: at 55°C and 60°C its single-source NMSE remains the lowest, while Physical-only, Data-NGO, Neumann NGO, and PINO-style grow substantially. The direct inverse baselines FNO and DeepONet stay near NMSE 1 and Top-1 25% (chance) at every temperature, indicating substantially degraded extrapolation performance under the evaluated distribution shift—as expected, because they directly map high-dimensional observations to sources without explicitly modeling the environment-dependent propagation operator.

Table 10: Synthetic OGW2 single-source recovery at 30 dB. Five-seed means use paired observations and matched recovery settings; Lunz-FAC uses a learned direction. Oracle provides a true-transfer reference, without assuming a known intact response. Full results appear in Table 11.

Method	50°C NMSE	50°C Top-1 (%)	55°C NMSE	55°C Top-1 (%)	60°C NMSE	60°C Top-1 (%)
Oracle transfer	0.02	100.0	0.02	100.0	0.02	100.0
Physical-only	13.29	51.5	24.85	50.4	40.64	48.3
PASL-U	**0.83**	**100.0**	**7.93**	**75.1**	23.03	43.0
Data-NGO	35.62	25.0	126.10	25.0	194.02	25.0
Neumann NGO	493.00	25.0	181.44	27.7	165.32	25.0
Lunz-FAC	2.02	99.9	9.14	72.2	**22.13**	**51.3**
PINO-style	1279.81	25.0	1103.48	25.0	1070.30	26.5
DeepONet	1.09	26.2	1.10	24.9	1.10	25.0
FNO	1.07	25.0	1.07	25.0	1.07	25.0

Notably, the baselines exhibit a separation between waveform fidelity and localization: Lunz-FAC still localizes most sources at 50°C (Top-1 about 99%), but its reconstructed waveforms are less

accurate than those of PASL-U (single-source NMSE 1.884, about six times higher). This shows that PASL-U's advantage stems from its two-stage environment encoding: Stage A learns bounded correction modes at calibration environments, and Stage B explicitly learns an environment-to-code predictor, allowing it to predict how the propagation operator changes once the environment leaves the calibration range, thereby retaining both waveform and location information.

Table 11: Synthetic guided-wave recovery at 30 dB, five-seed mean $\pm$ sample SD. Top-1 and Top-2 are percentages.

50°C				
Method	Single NMSE	Two NMSE	Top-1 (%)	Top-2 recall (%)
Physical-only	13.293 ± 0.423	6.497 ± 0.128	51.5 ± 1.5	66.7 ± 0.9
PASL-U	$\mathbf{0.313 \pm 0.169}$	$\mathbf{0.210 \pm 0.103}$	$\mathbf{100.0 \pm 0.0}$	$\mathbf{100.0 \pm 0.0}$
Data-NGO	40.144 ± 7.801	19.113 ± 3.527	25.0 ± 0.0	55.6 ± 6.1
Neumann NGO	484.207 ± 53.887	229.940 ± 21.379	25.0 ± 0.0	50.0 ± 0.0
PINO-style	1415.500 ± 204.425	671.919 ± 78.320	26.4 ± 3.1	50.0 ± 0.9
DeepONet	1.027 ± 0.009	1.014 ± 0.005	25.0 ± 0.0	50.0 ± 0.0
FNO	1.014 ± 0.001	1.007 ± 0.000	25.0 ± 0.0	50.0 ± 0.0
Lunz-FAC	1.884 ± 0.231	0.993 ± 0.108	99.1 ± 1.5	99.7 ± 0.5

55°C				
Method	Single NMSE	Two NMSE	Top-1 (%)	Top-2 recall (%)
Physical-only	24.848 ± 0.822	12.205 ± 0.272	50.4 ± 1.7	57.2 ± 0.5
PASL-U	$\mathbf{4.454 \pm 2.260}$	$\mathbf{2.313 \pm 1.198}$	$\mathbf{86.3 \pm 11.3}$	$\mathbf{86.8 \pm 12.1}$
Data-NGO	134.870 ± 11.086	65.359 ± 4.374	25.0 ± 0.0	53.3 ± 4.6
Neumann NGO	91.744 ± 37.723	43.962 ± 16.129	32.1 ± 9.3	50.1 ± 0.2
PINO-style	1278.157 ± 138.667	618.166 ± 55.851	25.0 ± 0.0	50.0 ± 0.0
DeepONet	1.029 ± 0.010	1.015 ± 0.005	25.0 ± 0.0	50.0 ± 0.0
FNO	1.014 ± 0.001	1.007 ± 0.001	25.0 ± 0.0	50.0 ± 0.0
Lunz-FAC	8.838 ± 0.801	4.406 ± 0.497	70.6 ± 14.5	77.6 ± 6.0

60°C				
Method	Single NMSE	Two NMSE	Top-1 (%)	Top-2 recall (%)
Physical-only	40.637 ± 1.051	20.071 ± 0.337	48.3 ± 1.8	52.3 ± 0.5
PASL-U	$\mathbf{16.662 \pm 6.846}$	$\mathbf{8.341 \pm 3.494}$	50.7 ± 3.5	$\mathbf{65.2 \pm 7.2}$
Data-NGO	202.335 ± 7.631	99.356 ± 5.704	25.0 ± 0.0	50.0 ± 0.0
Neumann NGO	165.572 ± 12.819	81.349 ± 7.496	25.0 ± 0.0	51.3 ± 2.8
PINO-style	1140.368 ± 70.743	558.278 ± 24.258	26.1 ± 2.5	50.0 ± 0.0
DeepONet	1.032 ± 0.012	1.015 ± 0.005	25.0 ± 0.0	50.0 ± 0.0
FNO	1.014 ± 0.001	1.006 ± 0.000	25.0 ± 0.0	50.0 ± 0.0
Lunz-FAC	21.401 ± 2.742	10.618 ± 1.494	$\mathbf{53.7 \pm 10.3}$	63.8 ± 6.5