Locality and global information in elliptic solvers

Mantis Discovery

Abstract

We study how much global information a local approximation to an elliptic inverse needs. At each grid point, the solver receives nearby right-hand-side values and a continuous encoding of the full input; both the encoder and reconstruction may be nonlinear. For the Dirichlet Laplacian on an n^d grid, uniform relative energy error at any fixed tolerance below one requires order $n^d/(w+1)^d$ shared coordinates, where w is the local radius. A linear construction attains this order. We give finite bounds, leading constants as the tolerance approaches one, and a tradeoff between cumulative local propagation and the total dimension of global messages. The lower bound combines a Borsuk–Ulam argument with canonical correlations of the Green matrix. For a prescribed linear measurement map, it also gives an explicit pair of inputs that the solver cannot distinguish. Experiments on the two-dimensional Poisson problem compare local convolutional, moment-based, Fourier, and multiscale models. A model using 16 sine moments has mean energy error 6.1% on smooth forcings and 67.0% on white forcings at the training resolution. Broader training reduces some of this gap, whereas none of the tested learned initial guesses reduces median end-to-end solve time at the largest timed grid.

1 Introduction

Learned elliptic solvers often combine local updates with global information. Fourier layers, integral operators, and multiscale architectures represent that information in different ways [9]. Their accuracy is usually evaluated on a prescribed distribution of coefficients or forcing functions. A solver required to work for every right-hand side faces a stronger problem, because important components of an input can be absent from the training distribution.

We fix the elliptic operator and allow each grid point to use uncompressed local input together with a shared continuous encoding. The local computation and the encoder can both be nonlinear. If the local radius is w, the grid has n^d points, and relative energy error is at most a fixed $\varepsilon < 1$ for every right-hand side, then the smallest encoding dimension satisfies

$$m_* = \Theta_{d,\varepsilon}\left(\frac{n^d}{(w+1)^d}\right), \qquad w = o(n). \tag{1}$$

The count is of real coordinates computed from each input, not trained parameters, bits, or arithmetic operations. Local linear approximations with a low-rank correction attain the same order. Thus the lower bound concerns information available to the reconstruction rather than its expressive power.

The finite bound uses forcing and observation sets that are separated beyond the local radius. Canonical correlations of the discrete Green matrix quantify the remaining dependence. Borsuk–Ulam then bounds the dimension of a continuous encoding that can preserve it. Polynomial approximation supplies a construction of the same order. Repeating local updates and global messages replaces the local radius by the cumulative propagation radius and the encoding dimension by the total message dimension. We also treat asymptotic constants, heterogeneous coefficients, and nonlinear solution maps near a regular point.

The experiments separate limits imposed by a representation from effects of training and changes in the forcing distribution. For encoders built from known linear measurements, a singular-value calculation constructs opposite inputs with identical local data and shared measurements. We compare these pairs with held-out errors, numerical input optimization, longer training, and broader forcing distributions. Conjugate-gradient and multigrid baselines then test whether the learned approximation reduces the cost of reaching a prescribed residual tolerance.

1.1 Related work

Nonlinear approximation theory studies continuous compression and reconstruction [5]. The sphere argument used here is a standard lower bound for manifold widths; Siegel [19, Proposition 1] gives a convenient formulation. Dahlke, Novak, and Sickel [3, 4] compare linear and nonlinear approximation of elliptic solution operators on regular input classes, and Franco, Manzoni, and Zunino [7] study latent dimension in reduced models of parametrized PDEs. Here the output also receives uncompressed local input. The main additional step is to remove that input on selected observation sets and calculate the dimension needed for the remaining elliptic response.

Low-rank Schur-complement corrections are used in positive definite preconditioners [1, 13]. Polynomial preconditioning and coarse corrections are established numerical methods [17, 21]. We use these methods only to construct an upper bound. The lower bound permits an arbitrary local matrix and, more generally, a continuous encoder with nonlinear reconstruction; it does not assume a factorization or a particular coarse solver.

Universality and information lower bounds use different quantifiers. Lanthaler, Li, and Stuart [11] prove operator approximation results with spatial averaging and variable channel dimension, while Lanthaler and Stuart [10] study parameter complexity for classes of target operators. Our target operator is fixed and the error condition ranges over every grid forcing. The propagation bounds of Tesán et al. [20] restrict local information flow; our result also counts global information. A sparse hierarchy can have large reach in the original grid, so it need not meet a small-radius assumption.

Structured changes in inputs and resolution have been used to evaluate neural operators [18]. Learned initial guesses for Krylov methods have also been studied [6]. Our experiments use these comparisons to interpret the information bound rather than to propose a Fourier architecture or a warm-start method. They concern one fixed Dirichlet problem, not a multi-equation benchmark.

2 Model and dimension bounds

Let $N = n^d$, with fixed $d \geq 1$, and index vectors by $[n]^d$, where $[n] = \{1, \dots, n\}$. The unscaled Dirichlet Laplacian and its inverse are

$$L_n = \operatorname{tridiag}(-1, 2, -1), \qquad A = A_{n,d} = \sum_{a=1}^{d} I_n^{\otimes(a-1)} \otimes L_n \otimes I_n^{\otimes(d-a)}, \qquad X = A^{-1}. \tag{2}$$

Write $\|u\|_A = (u^T A u)^{1/2}$. Restoring the physical factor $(n+1)^2$ rescales the inverse but leaves relative errors and condition numbers unchanged.

We first allow unrestricted transverse input and impose locality only in the first coordinate:

$$\mathcal{N}_w(i) = \{j \in [n]^d : |i_1 - j_1| \leq w\}, \qquad w \in \{0, 1, \dots\}. \tag{3}$$

A solver with m shared real coordinates has the form

$$\Phi_i(b) = \psi_i\big(b|_{\mathcal{N}_w(i)}, \mathcal{E}(b)\big), \qquad \mathcal{E} \in C(\mathbb{R}^N, \mathbb{R}^m). \tag{4}$$

The reconstruction functions ψ_i are unrestricted; all coefficients and trained parameters are fixed before b is supplied. Let $m_*(n, w, \varepsilon)$ be the smallest m for which

$$\|\Phi(b) - Xb\|_A \le \varepsilon \|Xb\|_A \quad \text{for every } b \in \mathbb{R}^N, \qquad 0 < \varepsilon < 1. \tag{5}$$

The space $\mathbb{R}^0$ is a singleton. Continuity of the encoder is essential: without regularity, a single real number can encode a finite-dimensional vector set-theoretically. No continuity of the reconstruction is needed for the lower bound.

For comparison, let $r_*(n, w, \Gamma)$ be the minimum rank of C among real symmetric decompositions

$$P = D + C \succ 0, \qquad D_{ij} = 0 \text{ if } |i_1 - j_1| > w, \qquad \kappa_A(P) \le \Gamma, \tag{6}$$

where

$$\kappa_A(P) = \frac{\lambda_{\max}(A^{1/2} P A^{1/2})}{\lambda_{\min}(A^{1/2} P A^{1/2})}. \tag{7}$$

This is the spectral condition number of a symmetric matrix, rather than the singular-value condition number of PA. Neither D nor C is required to be positive separately. The standard error identities for symmetric preconditioning [17] give

$$\sup_{b \ne 0} \frac{\|Pb - Xb\|_A}{\|Xb\|_A} = \|A^{1/2} P A^{1/2} - I\|_2, \qquad \min_{\gamma > 0} \|\gamma A^{1/2} P A^{1/2} - I\|_2 = \frac{\kappa_A(P) - 1}{\kappa_A(P) + 1}. \tag{8}$$

A factorization $C = UV^T$ needs only $V^T b \in \mathbb{R}^{\operatorname{rank} C}$ as shared information. Since scaling does not change support or rank,

$$m_*(n, w, \varepsilon) \le r_*(n, w, \Gamma), \qquad \Gamma = \frac{1 + \varepsilon}{1 - \varepsilon}. \tag{9}$$

To state the finite bounds, define

$$\lambda_j = 4 \sin^2 \frac{\pi j}{2(n+1)}, \qquad \mu_k = \sum_{a=1}^{d-1} \lambda_{k_a}, \qquad \theta_k = \operatorname{arcosh}(1 + \mu_k/2), \tag{10}$$

$$\ell = w + 1, \qquad p = \left\lfloor \frac{n + 1 - \ell}{2\ell} \right\rfloor, \qquad \xi = \tanh^2\!\left(\frac{\operatorname{arcosh}(1/\varepsilon)}{2\ell} \right). \tag{11}$$

Here $j \in [n]$, $k \in [n]^{d-1}$, and for $d = 1$ there is one empty transverse index with $\mu_k = \theta_k = 0$. When $3\ell \le n + 1$, set

$$R^-(n, w, \Gamma) = \#\left\{ (j, k) \in [p] \times [n]^{d-1} : \frac{\cos(\pi j/(2p+1))}{\cosh(\ell \theta_k)} > \varepsilon \right\}, \tag{12}$$

$$R^+(n, w, \Gamma) = \#\left\{ (j, k) \in [n] \times [n]^{d-1} : \lambda_j + (1 - \xi)\mu_k < 4\xi \right\}. \tag{13}$$

Throughout these counts, $\varepsilon = (\Gamma - 1)/(\Gamma + 1)$. All thresholds are strict; an equality mode already meets the target.

Theorem 2.1. *For $3(w + 1) \le n + 1$ and $0 < \varepsilon < 1$,*

$$R^- \le m_*(n, w, \varepsilon) \le r_*(n, w, \Gamma) \le R^+. \tag{14}$$

The upper bound has a linear realization $P = D + C$ with $D \succ 0$, $C \succeq 0$, $\operatorname{rank} C = R^+$, and relative energy error at most ε.

For graph-distance locality, replace (3) by

$$\mathcal{N}_w^{\mathrm{g}}(i) = \{j \in [n]^d : \textstyle\sum_a |i_a - j_a| \le w\}, \tag{15}$$

and denote the minima by m_*^{g} and r_*^{g}. The smaller neighborhoods can only increase the required dimension.

Theorem 2.2. *For fixed $d \ge 1$, fixed $0 < \varepsilon < 1$, and $\ell = w + 1 = o(n)$, all four minima satisfy*

$$m_*,\ r_*,\ m_*^{\mathrm{g}},\ r_*^{\mathrm{g}} = \Theta_{d,\varepsilon}((n/\ell)^d). \tag{16}$$

If $\Gamma \to \infty$ and $\ell\sqrt{\Gamma} = o(n)$, the one-coordinate minima m_* and r_* each equal

$$\left(\frac{\omega_d}{\pi^d} + o(1)\right) \frac{n^d}{\ell^d \Gamma^{d/2}}, \tag{17}$$

where ω_d is the Euclidean unit-ball volume. The rank conclusions also hold under separate positivity of D and C.

Sections 3 and 4 prove Theorem 2.1. The supplementary material counts the relevant modes and gives a graph-distance refinement with a different leading constant under the stronger condition $\ell\sqrt{\Gamma} = o(n^{1/d})$. The asymptotic statements with growing Γ describe the regime $\varepsilon \uparrow 1$; the finite estimates hold at any fixed tolerance.

3 Information across separated sets

Let A be any real symmetric positive definite matrix and write $X = A^{-1}$. For disjoint index sets I, J, define

$$T_{J,I} = X_{JJ}^{-1/2} X_{JI} X_{II}^{-1/2}. \tag{18}$$

The singular values $\sigma_1 \ge \cdots \ge \sigma_q \ge 0$, with $q = \min(|I|, |J|)$, are the canonical correlations between the two blocks of $X_{I\cup J, I\cup J}$ and are less than one. The Schur-complement variational principle [12, Theorem 2.2] gives

$$\min_{v:\, v_J = y} v^T A v = y^T X_{JJ}^{-1} y. \tag{19}$$

Thus a prescribed error on J has an unavoidable energy cost, even when the error outside J is chosen optimally.

Theorem 3.1. *Suppose $I \cap \mathcal{N}_w(j) = \varnothing$ for every $j \in J$. A solver of the form (4) satisfying (5) obeys*

$$m \ge \#\{a : \sigma_a(T_{J,I}) > \varepsilon\}. \tag{20}$$

The same conclusion holds when accuracy is required only for inputs supported on I with $b_I^T X_{II} b_I = 1$.

Proof. Assume $m < q$ and $\sigma_{m+1} > \varepsilon$. Let V be the span of the first $m+1$ right singular vectors of $T_{J,I}$. For v on the unit sphere of V, set

$$b(v)_I = X_{II}^{-1/2} v, \qquad b(v)_{I^c} = 0. \tag{21}$$

The solution $Xb(v)$ has unit energy norm. Apply the classical Borsuk–Ulam obstruction for continuous encodings [5, 19] to $v \mapsto \mathcal{E}(b(v))$. Some antipodal pair has the same encoding. Its local

inputs on J are both zero, so the computed outputs also agree there. Call their common restriction y.

Equation (19) and the assumed accuracy give

$$\|X_{JJ}^{-1/2} y - T_{J,I} v\|_2 \le \varepsilon, \qquad \|X_{JJ}^{-1/2} y + T_{J,I} v\|_2 \le \varepsilon.$$

The triangle inequality implies $\|T_{J,I} v\|_2 \le \varepsilon$, whereas $v \in V$ implies $\|T_{J,I} v\|_2 \ge \sigma_{m+1}$. This is a contradiction. If $m \ge q$, the conclusion is automatic. The argument includes $m = 0$, when the sphere consists of two points. □

3.1 An explicit test for fixed linear measurements

When the shared information is a function of known measurements Bb, opposite inputs can be found by linear algebra. Let Z have orthonormal columns spanning

$$\ker(B_{:,I} X_{II}^{-1/2}), \qquad \beta_{I,J}(B) = \|T_{J,I} Z\|_2.$$

If the kernel is zero, set $\beta_{I,J}(B) = 0$. Otherwise, a leading right singular vector v of $T_{J,I} Z$ gives $b_I = X_{II}^{-1/2} Z v$ and $b_{I^c} = 0$. Then $Bb = B(-b) = 0$ and both local inputs vanish on J. By (19), at least one of the two relative energy errors is at least $\beta_{I,J}(B)$. This is the usual nullspace calculation for a radius of information [22], applied after excluding local input. The singular-value min–max principle also gives $\beta_{I,J}(B) \ge \sigma_{m+1}(T_{J,I})$ when B has at most m rows and $m < q$ [8].

This test depends on the measurement map and the elliptic operator, not on the trained reconstruction. A nonlinear map applied to Bb cannot distinguish the pair. In numerical experiments, the nullspace residual and the difference between the two computed outputs must also be checked; a floating-point calculation is not an exact indistinguishability certificate.

3.2 Comparison of correlation counts

For a positive definite two-block matrix Z, let $\nu_\Gamma(Z)$ count the canonical correlations strictly larger than $(\Gamma - 1)/(\Gamma + 1)$. Write $J_0 = \mathrm{diag}(I, -I)$ and let $n_-(H)$ count the strictly negative eigenvalues of a Hermitian matrix H.

Lemma 3.2. *Canonical correlations are invariant under inversion and invertible block diagonal congruence. Moreover,*

$$\nu_\Gamma(Z) = n_-(\Gamma Z - J_0 Z J_0). \tag{22}$$

If $aZ \preceq Y \preceq bZ$, with $0 < a \le b$ and $\eta = b/a$, then

$$\nu_\Gamma(Y) \ge \nu_{\eta\Gamma}(Z). \tag{23}$$

If $\mathrm{rank}(Y - Z) \le s$, *then*

$$|\nu_\Gamma(Y) - \nu_\Gamma(Z)| \le 2s. \tag{24}$$

Proof. Use the standard block normalization and singular-value decomposition [1, 8]. In the resulting 2×2 blocks, inversion changes a normalized cross entry σ_j to $-\sigma_j$, and $\Gamma Z - J_0 Z J_0$ has eigenvalues $(\Gamma - 1) \pm (\Gamma + 1)\sigma_j$. Sylvester's law of inertia gives (22). The comparison follows from

$$\Gamma Y - J_0 Y J_0 \preceq a(\eta\Gamma Z - J_0 Z J_0)$$

and eigenvalue monotonicity. For (24), the two matrices in (22) differ by rank at most $2s$; the standard rank-perturbation inequality for Hermitian eigenvalue counts gives the result [8]. □

4 Finite bounds on the grid

4.1 Correlations between separated slices

Retain the slices with first coordinates

$$\ell, 2\ell, \dots, 2p\ell. \tag{25}$$

Let I contain the odd-numbered slices, let J contain the even-numbered slices, and put $E = I \cup J$. Their separation is at least $\ell > w$, so a forcing supported on I gives zero local input at every point of J. The inverse of X_{EE} is the Schur complement obtained by eliminating the unselected points. A transverse sine transform reduces this calculation to $L_n + \mu_k I$.

The inverse of a tridiagonal Dirichlet matrix has a hyperbolic representation [15]. At the sampled points in (25), it gives a tridiagonal Schur complement with off-diagonal entries $-c$ and diagonal entries a, apart from the final entry a_R:

$$c = \frac{\sinh\theta}{\sinh(\ell\theta)}, \qquad a = 2c\cosh(\ell\theta), \tag{26}$$

$$a_R = \sinh\theta\{\coth(\ell\theta) + \coth(h_R\theta)\}, \qquad h_R = n + 1 - 2p\ell \in [\ell, 3\ell). \tag{27}$$

Here $\theta = \operatorname{arcosh}(1 + \mu/2)$. At $\theta = 0$, the formulas mean $c = 1/\ell$, $a = 2/\ell$, and $a_R = 1/\ell + 1/h_R$. The final interval is retained in the formula; no divisibility assumption on $n + 1$ is needed. Since $h_R \geq \ell$, we have $a_R \leq a$.

Proposition 4.1. *The canonical correlations of X_{EE} in the partition (I, J) dominate, after decreasing rearrangement,*

$$\left\{ \frac{\cos(\pi j/(2p+1))}{\cosh(\ell\theta_k)} : j \in [p],\ k \in [n]^{d-1} \right\}. \tag{28}$$

Equality holds when $h_R = \ell$.

Proof. For one transverse mode and $h_R = \ell$, the two diagonal blocks of the Schur complement are aI_p, while the cross block is $-cB_p$, with B_p unit bidiagonal. Its singular values are $2\cos(\pi j/(2p+1))$, $j \in [p]$, by the discrete sine spectrum of the path matrix. Normalization therefore gives the displayed values because $c/a = 1/(2\cosh(\ell\theta))$.

For $h_R > \ell$, only the final diagonal entry in the even block decreases. The normalized cross block changes from T_0 to $T_0 R$, where R is diagonal with one entry $\sqrt{a/a_R} \geq 1$ and all other entries one. Thus $T_0 R^2 T_0^T \succeq T_0 T_0^T$, so no singular value decreases. The transverse transform acts separately within I and J; the full spectrum is the union of the mode spectra. Invariance under inversion in Lemma 3.2 transfers the result from the Schur complement to X_{EE}. □

Proposition 4.1 and Theorem 3.1 prove the lower bound in Theorem 2.1. The selected slices form one partition, so the encoding dimension is counted once. The sine transform only evaluates Green-matrix correlations; it imposes no symmetry or invariance assumption on the solver.

4.2 Polynomial approximation and spectral correction

For $0 < a < b$, the classical Chebyshev residual polynomial [17, Section 12.3.2] is

$$q(t) = \frac{T_\ell((b + a - 2t)/(b - a))}{T_\ell((b + a)/(b - a))}, \qquad p_w(t) = \frac{1 - q(t)}{t}, \quad w = \ell - 1. \tag{29}$$

The polynomial p_w has degree at most w. With $a/b = \xi$ from (11), the standard residual bounds are

$$|q(t)| \le \varepsilon \quad (a \le t \le b), \qquad \varepsilon < q(t) < 1 \quad (0 < t < a). \tag{30}$$

In particular, $p_w(t) > 0$ on $(0, b]$. We use this established approximation with intervals chosen separately for the transverse modes.

Upper bound in Theorem 2.1. In transverse sine coordinates, let $K_k = L_n + \mu_k I$ and choose the interval endpoints

$$b_k = 4 + \mu_k, \qquad a_k = \xi b_k.$$

Define D by the blocks $p_{w,k}(K_k)$, using (29) on each interval. These blocks are positive definite and have bandwidth at most w, since K_k is tridiagonal. Transforming back changes only transverse coordinates, so D has the required range in the first coordinate.

Let $v_{j,k}$ be the tensor sine eigenvectors of A and $t_{j,k} = \lambda_j + \mu_k$. Set

$$C = \sum_{t_{j,k} < a_k} \frac{q_k(t_{j,k}) - \varepsilon}{t_{j,k}} \, v_{j,k} v_{j,k}^T. \tag{31}$$

Each coefficient is positive. The inclusion condition is $\lambda_j + (1 - \xi)\mu_k < 4\xi$, hence $\operatorname{rank} C = R^+$. On a corrected mode, $A^{1/2}(D + C)A^{1/2}$ has eigenvalue $1 - \varepsilon$. Every other eigenvalue is $1 - q_k(t_{j,k}) \in [1 - \varepsilon, 1 + \varepsilon]$. This proves both the error bound and the condition-number bound. The shared coordinates are the inner products of b with the corrected eigenvectors. □

For graph-distance locality, use one interval with $b = 4d$ and $a = 4d\xi$. A degree-w polynomial in A connects only points at graph distance at most w. The same spectral correction gives

$$R^{\mathrm{g}} = \#\{\lambda \in \operatorname{spec}(A) : \lambda < 4d\xi\}, \qquad R^- \le m_*^{\mathrm{g}} \le r_*^{\mathrm{g}} \le R^{\mathrm{g}}. \tag{32}$$

Eigenvalues are counted with multiplicity. This construction bounds dimension, not the cost of computing its eigenvectors. Section 6 reports that cost separately for heterogeneous coefficients.

5 Local propagation and global layers

Consider an algorithm with L stages. Its initial state at i depends continuously on b_i and fixed data. At stage t, a continuous global operation computes $z_t \in \mathbb{R}^{m_t}$, and the next state at i depends continuously on these messages and on states within distance a_t of i. State dimensions and local computation are unrestricted. Distance can mean graph distance or $|i_1 - j_1|$. The final output is a pointwise function of the final state and the messages.

Proposition 5.1. *Put $W = \sum_{t=1}^{L} a_t$ and $M = \sum_{t=1}^{L} m_t$. The output has the form (4) with radius W and a continuous encoding into $\mathbb{R}^M$. Hence uniform relative error at most ε implies*

$$M \ge R^-(n, W, (1 + \varepsilon)/(1 - \varepsilon)) \qquad \text{if } 3(W + 1) \le n + 1. \tag{33}$$

Proof. Use the message sequence $(z_1(b), \dots, z_L(b))$ as the encoder. It is continuous. With that sequence fixed, induction on the stages shows that the state at i depends only on initial inputs within distance $\sum_{s=1}^{t} a_s$. The induction step is the triangle inequality for the chosen distance. Theorem 2.1 gives the bound. □

At fixed tolerance and $W = o(n)$, this yields

$$\left(\sum_{t=1}^{L} m_t\right)\left(1 + \sum_{t=1}^{L} a_t\right)^d \geq c_{d,\varepsilon} n^d. \tag{34}$$

For stage radius at most s and message dimension at most h, we obtain $Lh(1 + Ls)^d \geq c_{d,\varepsilon} n^d$ when $Ls = o(n)$. An algorithm without global messages therefore cannot achieve a fixed relative error with $W = o(n)$, regardless of local state dimension.

A Fourier layer retains a set K_t of complex coefficients from a c_t-channel state and combines the resulting global field with a local update [9, 14]. Its global dependence can be represented using at most $2c_t|K_t|$ real coordinates; conjugate symmetry may reduce this number. For pointwise local updates and lifting, (34) requires

$$2\sum_{t=1}^{L} c_t|K_t| \geq c_{d,\varepsilon} n^d. \tag{35}$$

A spatial average of c_t channels uses at most c_t coordinates. Nonlinearities may generate new frequencies between layers; the argument remains valid because it retains every layer's message. Local convolutions contribute their radii to W.

The count includes input-dependent pooling, normalization, and preprocessing. Coarse-grid transfers must be measured by their reach in the original grid, so a multilevel method may fall outside the small-W regime despite sparse updates. The quantity M is a real dimension, not a bit count or a communication cost; distributed Laplacian bounds in edge-capacity models use a different resource [2]. Equation (35) concerns uniform error rather than average error, and its conservative count can make the finite bound uninformative on small grids.

6 Empirical evaluation

We evaluate the fixed two-dimensional Dirichlet problem $A_h u = b$, with $A_h = (n+1)^2 A_{n,2}$, which separates variation in the forcing from variation in the coefficient field. The experiments compare average prediction error, pairs that a prescribed encoder cannot distinguish, and iterative-correction cost. The models are controlled implementations, not reproductions of a public benchmark.

6.1 Data, models, and error measures

Training uses 512 forcings and 128 validation forcings on a 31×31 grid. Smooth forcings have independent Gaussian sine coefficients divided by $k_1^2 + k_2^2$, whereas white forcings have undamped Gaussian coefficients. Each sample is scaled to unit root mean square amplitude during data generation; this defines the input distribution and is not part of the model. A sine-transform solve provides the discrete labels. Primary tests use 128 new forcings from each distribution at 31×31 and 63×63. Additional tests use band-limited and spatially localized forcings. White and band forcings at the finer grid contain new frequencies, so they change the input distribution as well as the resolution.

The local model has four five-point convolutional layers of width 16 and graph radius $w = 4$. The moment models add $m \in \{4, 16, 64\}$ fixed low-frequency sine measurements and a nonlinear decoder; their shared dimension is m. Fourier models use three spectral layers of width 12, with $k = 2$ or $k = 4$ modes in each retained direction, and pointwise updates [14]. Their conservative message dimensions are 576 and 2304. A four-level U-Net gives a multiscale comparison [16]; its

Model	Shared dimension	31×31 Smooth	31×31 White	63×63 Smooth	63×63 White
Local	0	0.599 ± 0.001	10.413 ± 0.709	0.857 ± 0.019	41.896 ± 2.923
Local + 4 moments	4	0.168 ± 0.003	1.831 ± 0.068	0.237 ± 0.009	7.321 ± 0.252
Local + 16 moments	16	0.061 ± 0.001	0.670 ± 0.030	0.095 ± 0.001	2.520 ± 0.083
Local + 64 moments	64	0.030 ± 0.001	0.903 ± 0.029	0.127 ± 0.007	3.651 ± 0.182
Fourier, $k = 2$	≤ 576	0.287 ± 0.002	4.399 ± 0.067	0.272 ± 0.001	17.235 ± 0.249
Fourier, $k = 4$	≤ 2304	0.262 ± 0.005	1.655 ± 0.146	0.246 ± 0.005	6.800 ± 0.575
U-Net	–	0.136 ± 0.025	1.201 ± 0.201	0.881 ± 0.088	7.172 ± 1.380

Table 1: Mean relative energy error after smooth-forcing training, with standard deviation across three seeds. Each mean uses 128 test inputs. Fourier dimensions are upper bounds on the full message sequence; the local and moment models have radius four. The U-Net is not assigned that radius. Errors are ratios, so values above one are possible.

coarse transfers have greater reach in the original grid. The supplementary material specifies the architectures and counts.

Every primary model is trained for 800 Adam updates using mean squared relative energy error, with the checkpoint selected by validation energy error. Three initialization and minibatch seeds use the same train, validation, and test sets. Reported deviations are standard deviations of the three seed means, not confidence intervals over independent data sets. The results also include relative L^2 error, residuals, quantiles, and individual values. Broader-training controls replace half the smooth training examples by white examples without increasing the sample count. The supplementary material gives the longer-training and larger-Fourier controls.

6.2 Prediction error under forcing and resolution changes

Increasing the number of sine moments improves accuracy on smooth forcings at the training grid, but not consistently on white forcings (Table 1). The 16-moment model has mean energy error 0.061 on smooth inputs and 0.670 on white inputs. With 64 moments, the respective errors are 0.030 and 0.903; at the finer grid, the white-input errors are 2.520 and 3.651. These are properties of the trained models, not of the best reconstruction allowed by each dimension.

The polynomial comparison in Figure 1 uses the same graph radius and correction rank, with an interval beginning at the first uncorrected eigenvalue. It is evaluated spectrally, so it isolates approximation error from the cost of implementation. The supplementary material reports its uniform residual bound. The learned models' errors after a distribution shift are therefore not determined by rank and radius alone.

Broader training improves white-input errors. With 16 moments, replacing half the training set by white forcings reduces the mean white error from 0.670 to 0.393 at the training grid and from 2.520 to 1.878 at the finer grid. For the primary Fourier model with $k = 4$, the same change reduces these errors from 1.655 to 0.560 and from 6.800 to 1.619. Longer optimization and a larger spectral layer do not remove the gap in these configurations. The controls identify a distributional effect; they do not evaluate every architecture or optimization method.

At larger resolutions, the finite lower bound is informative even with the conservative Fourier count. For the $k = 4$ architecture on a 255×255 grid, it gives a worst-case error of at least 0.802, while the observed mean error on smooth inputs is 0.266. The theorem concerns every input, whereas the supplementary table reports distributional averages. The larger white-input errors also reflect the changed forcing distribution and the fixed trained reconstruction.

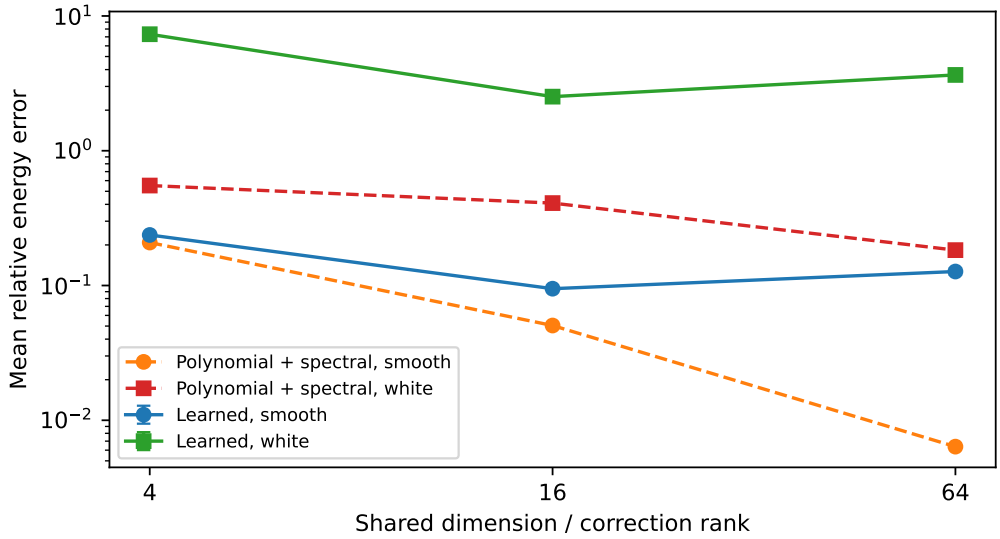


Figure 1: Mean energy error on the 63×63 grid as the number of sine moments or corrected eigenmodes increases. Solid curves are learned models trained on smooth inputs; error bars show one standard deviation across three seeds. Dashed curves are untrained degree-four Chebyshev approximations with exact correction of the first m eigenmodes. All curves use the same 128 test inputs per distribution. The plot shows average errors, not lower bounds on those averages.

6.3 Inputs with identical measurements

For the moment models, we use the construction in Section 3.1 with alternating rows separated by five grid steps. The forcing is supported on I, the local input is zero at each point of J, and every retained sine measurement vanishes. The construction uses the operator and measurement matrix, not the trained decoder. Table 2 reports the general dimension bound σ_{m+1} and the stronger value $\beta_{I,J}(B)$ for the prescribed measurements.

Every observed pair error exceeds its lower value. Relative measurement defects are below 1.44×10^{-18} in double precision, and relative output gaps on J are below 2.71×10^{-7} in the single-precision networks. At $n = 63, m = 16$, the prescribed-measurement value is 0.460 and the mean observed error is 1.321. Broader training leaves the measurements unchanged, so it cannot change this lower bound. Errors above $\beta_{I,J}(B)$ leave room to improve the reconstruction without increasing its information.

We also maximize relative energy error over inputs of unit root mean square amplitude using eight random restarts and 100 gradient steps per restart. An independent double-precision sine-transform solve checks the saved inputs. The supplementary material reports attained numerical errors, not certified maxima. Unlike the pairs above, these examples depend on the trained model and do not identify the Borsuk–Ulam witness for a general nonlinear encoder.

6.4 Iterative correction and computation cost

To compare complete solves, we use each learned output as an initial guess for conjugate gradient (CG), as in learned warm-start methods [6]. Every iterative method stops at relative Euclidean residual 10^{-6}, with no absolute tolerance. The baselines are zero-initialized CG and CG preconditioned by a geometric Galerkin V-cycle. A direct sine-transform solve is included because the fixed-coefficient problem admits that specialized algorithm. Warm-start time includes learned

n	m	σ_{m+1}	$\beta_{I,J}(B)$	Observed pair error
31	4	0.229	0.358	0.626 ± 0.040
31	16	0.017	0.092	0.555 ± 0.009
63	4	0.751	0.751	1.888 ± 0.347
63	16	0.385	0.460	1.321 ± 0.115
63	64	0.040	0.172	2.935 ± 0.147

Table 2: Errors on opposite inputs with zero retained sine measurements and zero local input on the observation rows. The first two error columns depend only on the operator and measurements. The last column is the larger error of the pair for each smooth-trained model, averaged over three seeds, with its standard deviation. The $n = 31, m = 64$ case is omitted because this selected forcing subspace has no measurement nullspace.

	Smooth forcing		White forcing	
Method	Iterations	Time (ms)	Iterations	Time (ms)
CG	334.8	27.78	320.9	29.16
MG-PCG	5.0	4.48	5.0	4.80
16 moments + CG	335.8	33.77	315.6	34.25
Fourier $k = 4$ + CG	337.5	45.07	324.6	42.80
Fourier $k = 4$, long + CG	336.6	43.21	330.6	42.39
U-Net + CG	367.1	42.67	362.5	42.33
U-Net, long + CG	367.8	42.58	363.4	41.55
DST	0.0	0.27	0.0	0.26

Table 3: Complete solves on the 127×127 grid. Iterations are means and times are medians over 32 new inputs per distribution. All iterative methods use the same residual threshold. Warm starts use initialization seed zero, fixed before timing; "long" denotes 3200 training updates instead of 800. Timing uses one CPU thread and batch size one. MG-PCG denotes multigrid-preconditioned CG; DST denotes the direct sine-transform solve, which has no iterations.

inference; training and preconditioner setup are reported separately.

None of the tested learned initial guesses reduces median complete-solve time on the 127×127 grid. On white forcings, the 16-moment model reduces the mean CG count from 320.9 to 315.6, but increases median time from 29.16 to 34.25 milliseconds once inference is included. The longer-trained Fourier and U-Net models do not reverse this result. Multigrid reaches the same tolerance in five iterations, and the direct transform is faster still because it exploits the constant-coefficient structure. Prediction accuracy alone therefore does not establish a computational benefit for these trained approximations.

For heterogeneous coefficients, we evaluate the analogous polynomial-and-correction construction as a preconditioner, not a learned map from coefficients to solutions. A degree-seven Chebyshev polynomial is supplemented by computed low eigenvectors with target $\varepsilon = 0.5$. The smooth and checkerboard cases have upper contrast bounds of 7 and 50. The correction keeps the preconditioned iteration count near 11, while its rank and setup cost grow with contrast. At $n = 63$ and contrast 50, it has rank 489 and takes 18.45 seconds to construct. Its median solve time is 9.71 milliseconds, compared with 4.77 milliseconds for multigrid, whose setup takes 2.35 milliseconds. The iteration reduction must therefore be considered with construction and application costs.

7 Discussion

The theory gives the necessary shared dimension for uniform elliptic inversion when each output also receives local input. The nonlinear lower bound and linear upper construction agree in order and have matching constants in the stated regime as $\varepsilon \uparrow 1$. The experiments show why this requirement differs from average test accuracy. A model can fit smooth forcings while a fixed measurement map provably misses other inputs; broader training can reduce observed error without changing the dimension bound.

The empirical results concern a fixed two-dimensional Poisson operator and CPU execution. The coefficient experiment evaluates classical preconditioners, not a neural operator trained over coefficient fields. They are neither an architecture ranking nor a general timing claim. The fixed-tolerance constants remain separated, the graph-distance refinement has an additional boundary condition, and average-case guarantees require a different analysis.

References

[1] Emmanuel Agullo, Eric Darve, Luc Giraud, and Yuval Harness. Low-rank factorizations in data sparse hierarchical algorithms for preconditioning symmetric positive definite matrices. *SIAM Journal on Matrix Analysis and Applications*, 39(4):1701–1725, 2018. doi: 10.1137/17M1151158.

[2] Ioannis Anagnostides, Christoph Lenzen, Bernhard Haeupler, Goran Zuzic, and Themis Gouleakis. Almost universally optimal distributed Laplacian solvers via low-congestion shortcuts. *Distributed Computing*, 36(4):475–499, 2023. doi: 10.1007/s00446-023-00454-0.

[3] Stephan Dahlke, Erich Novak, and Winfried Sickel. Optimal approximation of elliptic problems by linear and nonlinear mappings I. *Journal of Complexity*, 22(1):29–49, 2006. doi: 10.1016/j.jco.2005.06.005.

[4] Stephan Dahlke, Erich Novak, and Winfried Sickel. Optimal approximation of elliptic problems by linear and nonlinear mappings II. *Journal of Complexity*, 22(4):549–603, 2006. doi: 10.1016/j.jco.2006.04.001.

[5] Ronald A. DeVore, Ralph Howard, and Charles Micchelli. Optimal nonlinear approximation. *Manuscripta Mathematica*, 63(4):469–478, 1989. doi: 10.1007/BF01171759.

[6] Mohammad Sadegh Eshaghi, Cosmin Anitescu, Navid Valizadeh, Yizheng Wang, Xiaoying Zhuang, and Timon Rabczuk. NOWS: Neural operator warm starts for accelerating iterative solvers. *Computer Methods in Applied Mechanics and Engineering*, 458:118989, 2026. doi: 10.1016/j.cma.2026.118989.

[7] Nicola R. Franco, Andrea Manzoni, and Paolo Zunino. A deep learning approach to reduced order modelling of parameter dependent partial differential equations. *Mathematics of Computation*, 92(340):483–524, 2023. doi: 10.1090/mcom/3781.

[8] Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, second edition, 2013. doi: 10.1017/CBO9781139020411.

[9] 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. URL `https://jmlr.org/papers/v24/21-1524.html`.

[10] Samuel Lanthaler and Andrew M. Stuart. The parametric complexity of operator learning, 2025. URL https://arxiv.org/abs/2306.15924v4. arXiv:2306.15924v4.

[11] Samuel Lanthaler, Zongyi Li, and Andrew M. Stuart. Nonlocality and nonlinearity implies universality in operator learning. *Constructive Approximation*, 62(2):261–303, 2025. doi: 10.1007/s00365-025-09718-3.

[12] Chi-Kwong Li and Roy Mathias. Extremal characterizations of the Schur complement and resulting inequalities. *SIAM Review*, 42(2):233–246, 2000. doi: 10.1137/S0036144599337290.

[13] Ruipeng Li, Yuanzhe Xi, and Yousef Saad. Schur complement-based domain decomposition preconditioners with low-rank corrections. *Numerical Linear Algebra with Applications*, 23(4): 706–729, 2016. doi: 10.1002/nla.2051.

[14] 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. URL https://openreview.net/forum?id=c8P9NQVtmnO.

[15] Gérard Meurant. A review on the inverse of symmetric tridiagonal and block tridiagonal matrices. *SIAM Journal on Matrix Analysis and Applications*, 13(3):707–728, 1992. doi: 10.1137/0613045.

[16] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, volume 9351 of *Lecture Notes in Computer Science*, pages 234–241. Springer, 2015. doi: 10.1007/978-3-319-24574-4_28.

[17] Yousef Saad. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, Philadelphia, second edition, 2003. doi: 10.1137/1.9780898718003.

[18] Lennon Shikhman. Diagnosing failure modes of neural operators across diverse PDE families. *Transactions on Machine Learning Research*, 2026. URL https://openreview.net/forum?id=0S1LWZHQYn.

[19] Jonathan W. Siegel. Sharp lower bounds on the manifold widths of Sobolev and Besov spaces. *Journal of Complexity*, 85:101884, 2024. doi: 10.1016/j.jco.2024.101884.

[20] Lucas Tesán, Mikel M. Iparraguirre, David González, Pedro Martins, and Elías Cueto. On the under-reaching phenomenon in message passing neural PDE solvers: Revisiting the CFL condition. *Computer Methods in Applied Mechanics and Engineering*, 449:118476, 2026. doi: 10.1016/j.cma.2025.118476.

[21] Andrea Toselli and Olof Widlund. *Domain Decomposition Methods—Algorithms and Theory*, volume 34 of *Springer Series in Computational Mathematics*. Springer, 2005. doi: 10.1007/b137868.

[22] Joseph F. Traub, Grzegorz W. Wasilkowski, and Henryk Woźniakowski. *Information-Based Complexity*. Academic Press, Boston, 1988. ISBN 9780126975451.