WHAT SURVIVES HIGHER-ORDER SIGNAL REDUCTION? HARMONIC UNOBSERVABILITY AT VERTICES

Anonymous authors
Paper under double-blind review

ABSTRACT

Higher-order models represent signals across vertices, edges, faces, and higher-rank cells, yet collapsing these representations to a lower rank raises a fundamental question: which higher-order information remains observable after reduction? We characterize this information pathway for rank-shared boundary-mediated reductions and show that cross-rank transfer necessarily factors through boundary operators. Under the Hodge decomposition, the retained lower-rank representation can access boundary-visible gradient components, whereas harmonic components are structurally unobservable. We further show that this loss is caused by the observation interface itself rather than by the numerical treatment of singular higher-rank blocks. Motivated by this characterization, we introduce HARP, a HArmonic Readout Pipeline that explicitly exposes the missing harmonic information. HARP constructs a topology-preserving Morse complex, computes its harmonic basis, lifts the resulting representatives through a chain-compatible map, and projects them onto the harmonic subspace of the original complex. The resulting basis yields compact, sample-dependent harmonic coordinates that complement the information accessible through lower-rank reduction. Experiments on synthetic and real higher-order signals validate the predicted information pathway and show that HARP recovers useful information specifically when harmonic components are task-relevant. Together, these results identify a precise blind spot of rank-shared higher-order reduction and provide a compact interface for preserving the information hidden beyond it. Code will be released upon acceptance.

1 INTRODUCTION

Graphs provide a flexible representation for relational data and have become a standard modeling tool across a wide range of fields. However, their basic structure consists of vertices connected by edges, and therefore represents interactions primarily through pairwise relations. Many real systems instead exhibit *higher-order* interactions among groups of more than two entities, which cannot always be captured by independent pairwise connections. Topological deep learning (TDL) addresses this limitation by extending graphs to structures such as simplicial and cellular complexes, where signals can reside not only on vertices and edges but also on higher-rank elements such as triangles and faces (Bodnar et al., 2021a;b; Hajij et al., 2022; Papamarkou et al., 2024).

This richer representation, however, introduces a new challenge: information is distributed across multiple ranks and must eventually be combined for downstream tasks. Existing higher-order architectures typically process rank-specific signals through Hodge operators, message passing, or attention, and then exchange or aggregate information across ranks (Barbarossa & Sardellitti, 2020; Bodnar et al., 2021b; Hajij et al., 2022). For vertex-level tasks, this creates a fusion problem because higher-rank information must be transferred back to the vertex domain. Such fusion often relies on repeated cross-rank message passing, predefined aggregation schemes, or architecture-specific design choices (Papillon et al., 2023). Recent effective-operator approaches provide a principled alternative by eliminating higher-rank variables and collapsing their influence into a single vertex-level operator. In particular, Collapsed Effective Operators (CEO) (Krahn et al., 2026) use Schur complementation to integrate a graded higher-order system into a vertex-level representation, avoiding explicit multi-rank fusion while capturing topology-mediated interactions.

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.

Yet collapsing a higher-order representation raises a more fundamental question: *which higher-order information can actually survive the reduction?* A reduced operator can encode the influence of eliminated variables, but this does not imply that every component of a higher-order signal remains observable from the retained vertex space. Moreover, the resulting vertex-level operator can be dense, making it desirable to represent the information missed by reduction without retaining the full higher-order state. Rather than proposing another cross-rank aggregation rule, we study the information pathway itself and characterize what boundary-mediated reduction can and cannot observe.

We show that, within the rank-shared boundary-mediated operator class, cross-rank information can reach the vertex space only through the boundary operator. Under the Hodge decomposition (defined in Sec. 2), this interface transmits only the gradient component; in particular, harmonic components are structurally unobservable. Importantly, this loss occurs at the observation interface, not because of how the singular higher-rank block is handled. We further characterize admissible cross-rank couplings and show that this unobservability persists across the rank-shared class.

Motivated by this characterization, we introduce **HARP**, a **HArmonic Readout Pipeline** that explicitly extracts the harmonic information hidden from lower-rank reduction. HARP uses Forman discrete Morse theory (DMT) (Forman, 1998) to construct a smaller Morse complex while preserving the homology of the original complex. It then computes harmonic representatives in the reduced complex, lifts them through a chain-compatible map, and projects the lifted cycles onto the harmonic subspace of the original complex. The resulting basis yields compact, sample-dependent harmonic coordinates that complement rather than replace the reduced vertex representation. This provides a topology-preserving interface between higher-order signals and conventional lower-rank predictors without requiring a particular downstream architecture.

Our contributions are summarized as follows: **1)** We characterize the information pathway of rank-shared boundary-mediated reduction and prove that harmonic components are structurally unobservable from the retained lower-rank space, **2)** We introduce HARP, which uses topology-preserving Morse reduction, chain-compatible lifting, and harmonic projection to construct compact signal-dependent coordinates for the harmonic information, and **3)** We validate the predicted information pathway on synthetic and real higher-order signals, showing that HARP recovers useful information when harmonic information is present and consistently complements diverse lower-rank backbones.

2 PRELIMINARIES

Chain complexes and higher-order signals. Graphs represent pairwise interactions through vertices and edges, whereas higher-order structures extend them to multi-way interactions using elements of different ranks. A common representation is a simplicial complex, whose rank-k elements correspond to k-dimensional simplices, e.g., vertices, edges, and triangles at ranks 0, 1, and 2 (Hatcher, 2002).

Let $C_k \simeq \mathbb{R}^{n_k}$ denote the space spanned by the n_k oriented rank-k elements. A rank-k signal is represented as a vector in C_k. The boundary map $B_k \in \mathbb{R}^{n_{k-1} \times n_k}$ records the signed incidence between rank-k and rank-$(k-1)$ elements. The spaces $\{C_k\}_{k=0}^{K}$ and boundary maps $\{B_k\}_{k=1}^{K}$ form a *chain complex* when $B_{k-1}B_k = 0$, which states that the boundary of a boundary vanishes. Our analysis is formulated on finite chain complexes and does not rely on simplex-specific geometry.

Hodge decomposition. The *Hodge decomposition* expresses a rank-k signal as the orthogonal sum of gradient, curl, and harmonic components (Barbarossa & Sardellitti, 2020; Lim, 2020):

$$C_k = \underbrace{\operatorname{im} B_k^\top}_{\mathcal{G}_k\,:\,\text{gradient}} \oplus \underbrace{\operatorname{im} B_{k+1}}_{\mathcal{C}_k\,:\,\text{curl}} \oplus \underbrace{\left(\ker B_k \cap \ker B_{k+1}^\top\right)}_{\mathcal{H}_k\,:\,\text{harmonic}}. \tag{1}$$

The gradient component $\mathcal{G}_k$ is induced from the rank below, while the curl component $\mathcal{C}_k$ describes local circulations induced by cells one rank above. The harmonic subspace $\mathcal{H}_k$ captures the remaining global circulations that are neither gradients nor boundaries of higher-dimensional cells. Equivalently, $\mathcal{H}_k$ is the kernel of the k-th Hodge Laplacian Δ_k:

$$\Delta_k = B_k^\top B_k + B_{k+1}B_{k+1}^\top, \qquad \mathcal{H}_k = \ker \Delta_k. \tag{2}$$

Its dimension $\beta_k = \dim \mathcal{H}_k$ is the k-th Betti number, which counts the number of independent topological features at rank k, such as connected components $(k = 0)$ and loops $(k = 1)$. While β_k

is fixed by the underlying complex, the projection of a signal onto $\mathcal{H}_k$ varies across inputs. Our work focuses on this signal-dependent harmonic variation.

Higher-order reduction. Signals across all ranks can be represented jointly in the graded space $C_* = \bigoplus_{k=0}^{K} C_k$. A higher-order system can be reduced to a lower rank by eliminating its higher-rank variables via Schur complementation (Dorfler & Bullo, 2012). For rank-0 reduction, we partition C_* into the vertex space C_0 and the higher-rank space $\bigoplus_{k=1}^{K} C_k$, and write a symmetric operator $\mathcal{M}$ as:

$$\mathcal{M} = \begin{pmatrix} P & Q \\ Q^\top & R \end{pmatrix}, \qquad \mathcal{M}_\varepsilon^{\downarrow} = P - Q(R + \varepsilon I)^{-1} Q^\top, \quad \varepsilon > 0, \tag{3}$$

where P acts on the retained vertex space, R governs the eliminated higher-rank variables, and Q couples the two. The regularized Schur complement $M_\varepsilon^{\downarrow}$ defines an effective vertex-level operator that incorporates the influence of the eliminated higher-rank block through Q. To distinguish operator reduction from the transfer of an actual higher-rank signal, consider the block system with vertex source $f \in C_0$ and higher-rank source $z \in \bigoplus_{k=1}^{K} C_k$. Eliminating the higher-rank variable gives:

$$\mathcal{M}_\varepsilon^{\downarrow} x = f - T_\varepsilon z, \qquad T_\varepsilon := Q(R + \varepsilon I)^{-1}. \tag{4}$$

Thus, $\mathcal{M}_\varepsilon^{\downarrow}$ characterizes the reduced operator, whereas $T_\varepsilon z$ characterizes which components of a higher-rank input influence the retained representation. Collapsed Effective Operators (CEO) (Krahn et al., 2026) instantiate this construction using a graded higher-order operator and collapse its higher-rank blocks onto the vertex space. Our analysis concerns the homological realization in which the cross-rank maps are signed boundary operators satisfying the chain condition, i.e., $B_{k-1}B_k = 0$.

3 METHOD

We first characterize which higher-order signal components can survive boundary-mediated rank reduction. Based on this characterization, we then introduce HARP, a Harmonic Readout Pipeline that explicitly exposes the signal components that become unobservable after reduction.

3.1 BOUNDARY-MEDIATED OPERATOR CLASS

Consider a symmetric operator $\mathcal{M}$ whose within-rank interactions are governed by the Hodge Laplacians Δ_k and whose interactions between adjacent ranks are mediated by boundary maps B_k:

$$\mathcal{M} = \begin{pmatrix} \Delta_0 & -\gamma_1 B_1 & 0 & \cdots & 0 \\ -\gamma_1 B_1^\top & \Delta_1 & -\gamma_2 B_2 & \ddots & \vdots \\ 0 & -\gamma_2 B_2^\top & \Delta_2 & \ddots & 0 \\ \vdots & \ddots & \ddots & \ddots & -\gamma_K B_K \\ 0 & \cdots & 0 & -\gamma_K B_K^\top & \Delta_K \end{pmatrix}, \tag{5}$$

where $\gamma_k \geq 0$ is a coupling strength shared across cells of rank k, which we call a *rank-shared boundary coupling*. For the rank-0 reduction defined in equation 3, the cross-rank block is

$$Q = \begin{bmatrix} -\gamma_1 B_1 & 0 & \cdots & 0 \end{bmatrix}. \tag{6}$$

Hence, regardless of interactions within the eliminated block R, their influence on the vertex space must pass through B_1. Our main analysis considers this rank-shared reduction class. We separately examine how cell-dependent coupling weights affect the observable subspace in Appendix A.4.

3.2 HARMONIC UNOBSERVABILITY

The boundary-mediated interface in equation 6 restricts higher-rank information reaching the vertex space to that passing through B_1. Let $G_\varepsilon : C_1 \to C_1$ denote the rank-1 diagonal block of $(R+\varepsilon I)^{-1}$. Substituting equation 6 into the Schur complement (equation 3) yields the following factorization:

$$\mathcal{M}_\varepsilon^{\downarrow} = \Delta_0 - \gamma_1^2 B_1 G_\varepsilon B_1^\top. \tag{7}$$

Let $\Pi_{\mathcal{G}_1}$ denote the orthogonal projector onto the gradient subspace $\mathcal{G}_1 = \operatorname{im} B_1^\top$. Since the Hodge decomposition gives $\ker B_1 = \mathcal{C}_1 \oplus \mathcal{H}_1$, we have $B_1(I - \Pi_{\mathcal{G}_1}) = 0$, and therefore

$$M_\epsilon^\downarrow = \Delta_0 - \gamma_1^2 B_1 \left(\Pi_{\mathcal{G}_1} G_\epsilon \Pi_{\mathcal{G}_1}\right) B_1^\top. \tag{8}$$

The detailed derivation is provided in Appendix A.2.

Theorem 3.1 (Gradient Factorization). *For the operator class in equation 5, the rank-0 Schur correction factors through B_1 as in equation 7. Moreover, the induced vertex-level source $T_\varepsilon z$ depends on the rank-1 component z_1 only through its boundary response $B_1 z_1$. Hence, perturbations of z_1 within $\ker B_1 = \mathcal{C}_1 \oplus \mathcal{H}_1$ are unobservable from the retained vertex space.*

Theorem 3.1 establishes that the retained vertex space accesses rank-1 information only through B_1. For a curl signal $c = B_2 u$, the chain condition gives $B_1 c = B_1 B_2 u = 0$, while any harmonic signal $h \in \mathcal{H}_1$ also satisfies $B_1 h = 0$. Nevertheless, their structural roles differ: curl variation may remain visible to the rank-1/rank-2 interface through $B_2^\top$, whereas a harmonic signal $h \in \mathcal{H}_k$ satisfies both $B_k h = 0$ and $B_{k+1}^\top h = 0$ and is therefore annihilated by both adjacent boundary interfaces.

One may ask whether this limitation is specific to choosing B_1 as the cross-rank coupling. To address this question, we characterize boundary-compatible linear couplings between rank-1 and rank-0 spaces. These couplings preserve the signed-incidence locality and orientation structure of the original boundary interaction. The formal conditions are provided in Appendix A.3.

Theorem 3.2 (Coupling Rigidity). *Let $Y : C_1 \to C_0$ be a linear coupling satisfying: (i) signed-incidence locality, where each edge contribution follows its signed incidence pattern and may depend on intrinsic edge attributes, (ii) orientation covariance under edge reversal, and (iii) relabeling equivariance under vertex permutations. Then,*

$$Y = B_1 \Gamma, \tag{9}$$

where Γ is a diagonal matrix containing edge-wise coupling weights. If the coupling is independent of edge attributes and isotropic across edges, then $\Gamma = \gamma I$.

Theorem 3.2 shows that these structural conditions reduce the freedom in the coupling to edge-wise reweighting of the boundary operator. Following the isotropic rank-wise parameterization used in Krahn et al. (2026), our main analysis considers $\Gamma = \gamma I$. In this setting,

$$Y = \gamma B_1, \qquad \ker(Y) = \ker(B_1), \tag{10}$$

for $\gamma \neq 0$. Hence, the reference curl and harmonic directions remain unobservable. We analyze cell-dependent anisotropic reweighting in Appendix A.4, showing that the reweighted boundary operator remains blind to the harmonic subspace induced by its own Hodge decomposition. Consequently, anisotropic reweighting does not provide a lossless representation of the full higher-order signal.

A remaining question is whether the loss of harmonic information is caused by the singularity of R rather than by the observation interface. The harmonic subspaces of the eliminated ranks are contained in $\ker R$:

$$\mathcal{K}_\mathcal{H} := \bigoplus_{k=1}^{K} \mathcal{H}_k \subseteq \ker R. \tag{11}$$

Let $U_\mathcal{H} = [u_1, \ldots, u_\beta]$ be an orthonormal basis of $\mathcal{K}_\mathcal{H}$, where $\beta = \sum_{k=1}^K \beta_k$. Because the rank-1 component of every u_j lies in $\mathcal{H}_1 \subseteq \ker B_1$, the boundary-mediated interface satisfies:

$$QU_\mathcal{H} = 0. \tag{12}$$

Standard treatments of this singularity include the Moore–Penrose pseudoinverse (Moore, 1920; Penrose, 1955), Tikhonov regularization (Tikhonov, 1963), and assigning a finite positive mass to the harmonic directions. Although these approaches modify the treatment of the harmonic kernel, they do not modify the observation interface Q (details in Appendix A.5). Since $QU_\mathcal{H} = 0$, every contribution supported on $\mathcal{K}_\mathcal{H}$ vanishes after propagation to the retained vertex space.

Theorem 3.3 (Kernel Reduction). *For the rank-shared boundary-mediated coupling $Q = [-\gamma_1 B_1 \mid 0 \mid \cdots \mid 0]$, the harmonic subspace satisfies*

$$\mathcal{K}_\mathcal{H} \subseteq \ker R \qquad \textit{and} \qquad QU_\mathcal{H} = 0.$$

Consequently, the harmonic contribution to the reduced operator vanishes under the Moore–Penrose pseudoinverse, Tikhonov regularization, and any finite positive mass supported on $\mathcal{K}_\mathcal{H}$.

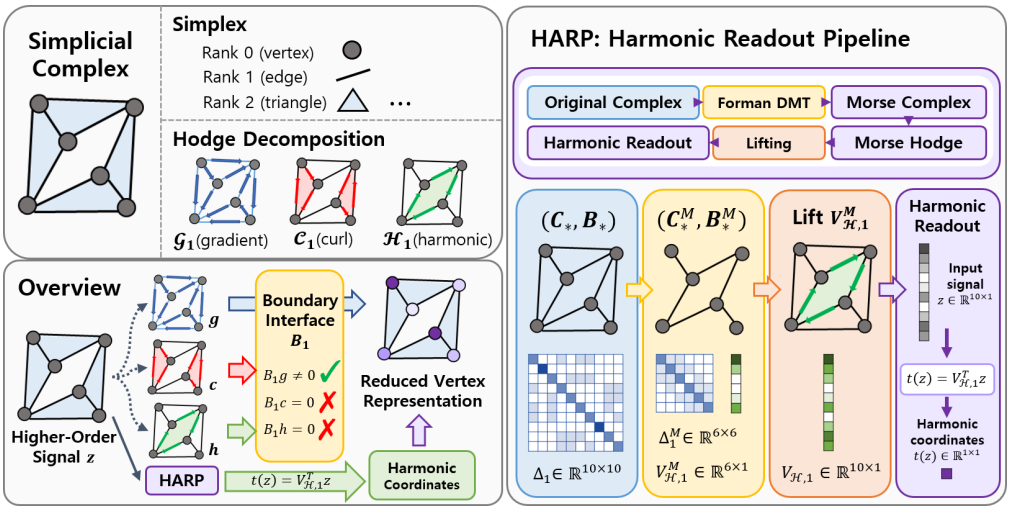

Figure 1: Overview of harmonic unobservability and HARP using a rank-2 simplicial complex with $\beta_1 = 1$. Left: for a rank-1 signal $z = g + c + h$, the boundary interface B_1 transmits gradient variation but annihilates curl and harmonic components, while HARP preserves the missing harmonic information through explicit coordinates. Right: Forman DMT constructs a topology-preserving Morse complex, on which HARP computes a harmonic basis, lifts and projects it in the original chain space, and obtains the readout $t(z) = V_{\mathcal{H},1}^\top z$.

Theorem 3.3 shows that harmonic information is lost at the observation interface, not through the numerical treatment of the singular higher-rank block. This structural limitation motivates HARP, which explicitly extracts harmonic coordinates rather than recovering them via the reduction operator.

3.3 HARP: Harmonic Readout Pipeline

The preceding analysis identifies harmonic components as a structurally hidden signal subspace under boundary-mediated reduction. As illustrated in Figure 1, HARP explicitly constructs this subspace and extracts signal-dependent harmonic coordinates. The pipeline consists of three stages: topology-preserving Morse reduction, harmonic basis construction, and harmonic readout.

Topology-preserving Morse reduction. Direct Hodge decomposition identifies the harmonic subspace directly on the original complex. To obtain a smaller topology-preserving chain complex together with an explicit correspondence to the original chain space, we apply Forman discrete Morse theory (DMT) (Forman, 1998). DMT constructs a smaller Morse complex (C_*^M, B_*^M) that is chain-homotopy equivalent to the original complex. DMT pairs incident cells of adjacent ranks through an acyclic matching and retains unmatched cells as critical cells. We keep all rank-0 cells unmatched to preserve the original vertex domain. Because a critical-cell coordinate is defined in the reduced Morse complex rather than the chain space, HARP uses a chain-compatible lifting map $E_k : C_k^M \to C_k$ satisfying

$$B_k E_k = E_{k-1} B_k^M. \tag{13}$$

Hence, every cycle in the Morse complex is lifted to a cycle in the original complex (Appendix B.2).

Harmonic basis construction. For each relevant rank k, we compute an orthonormal basis $V_{\mathcal{H},k}^M$ of the Morse harmonic subspace, i.e., $\ker \Delta_k^M$. Specifically, we construct the Morse Hodge Laplacian:

$$\Delta_k^M = (B_k^M)^\top B_k^M + B_{k+1}^M (B_{k+1}^M)^\top. \tag{14}$$

Because DMT preserves homology, $\dim \ker \Delta_k^M = \beta_k$, matching the harmonic dimension of the original complex. We then lift the Morse harmonic basis through E_k. The lifted vectors remain cycles because the lifting map is chain-compatible, i.e., $B_k E_k V_{\mathcal{H},k}^M = 0$. However, being a cycle does not necessarily imply being harmonic, since the cycle space decomposes as $\ker B_k = \operatorname{im} B_{k+1} \oplus \mathcal{H}_k$. We remove this boundary component and orthonormalize the remaining harmonic representatives:

$$V_{\mathcal{H},k} = \operatorname{orth}\left[\left(I - B_{k+1} B_{k+1}^\dagger\right) E_k V_{\mathcal{H},k}^M\right]. \tag{15}$$

Here, $B_{k+1}^{\dagger}$ denotes the Moore–Penrose pseudoinverse, making $B_{k+1}B_{k+1}^{\dagger}$ the orthogonal projector onto $\mathrm{im}\, B_{k+1}$. Thus, $I - B_{k+1}B_{k+1}^{\dagger}$ removes the boundary component, while $\mathrm{orth}(\cdot)$ orthonormalizes the remaining vectors. The resulting basis satisfies $B_k V_{\mathcal{H},k} = 0$ and $B_{k+1}^{\top} V_{\mathcal{H},k} = 0$, and hence spans the harmonic subspace $\mathcal{H}_k$ (Appendix B.3).

Signal-dependent harmonic readout. Collecting the rank-wise harmonic bases gives $V_{\mathcal{H}} = \mathrm{blkdiag}(V_{\mathcal{H},1}, \ldots, V_{\mathcal{H},K})$. For a higher-order signal $z = [z_1^{\top}, \ldots, z_K^{\top}]^{\top}$, HARP extracts:

$$t(z) = V_{\mathcal{H}}^{\top} z. \tag{16}$$

The vector $t(z)$ contains coefficients along harmonic directions hidden from the reduction interface.

Importantly, this readout is distinct from the Betti numbers. The Betti number β_k specifies only the dimension of $\mathcal{H}_k$ and is fixed for a given complex, whereas $t(z)$ varies across input signals on the same complex. For a multi-channel signal Z, the same operation is applied as $T(Z) = V_{\mathcal{H}}^{\top} Z$. Direct Hodge decomposition recovers the same harmonic subspace and serves as a baseline, while DMT provides a topology-preserving reduced complex together with an explicit chain-compatible correspondence to the original complex.

3.4 Using Harmonic Readouts

The harmonic readout defined above can be used either to augment a lower-rank representation or as a standalone higher-order representation. For vertex-level prediction, we broadcast the layer-wise harmonic coordinates $T^{(\ell)}$ to the vertex domain and concatenate them with the vertex representation:

$$\widetilde{H}^{(\ell)} = H^{(\ell)} \,\|\, \mathbf{1}_n \,\mathrm{vec}\Big(T^{(\ell)}\Big)^{\top}. \tag{17}$$

This injection is deliberately architecture-agnostic: HARP specifies the harmonic information to expose, while the downstream model determines how to use it. For tasks defined directly on the higher-order signal, $t(z)$ can instead be used as a readout-only representation without vertex features.

4 Experiments

Our experiments are designed to evaluate the structural predictions of Sec. 3, rather than to position HARP as a generic performance booster for graph learning. We aim to answer three questions: **1)** does rank-shared reduction remove harmonic information as predicted? **2)** does explicitly preserving the missing harmonic coordinates provide useful and compact information on real higher-order signals? **3)** are the observed gains specific to the certified harmonic subspace rather than topology alone or arbitrary low-dimensional projections?

4.1 Experimental Setup

Datasets and tasks. We evaluate HARP on synthetic chain complexes and real higher-order signals. The synthetic dataset contains an annulus ($\beta_1 = 1$), a two-hole strip ($\beta_1 = 2$), and the boundary of a tetrahedron ($\beta_2 = 1$), allowing the harmonic component to be controlled exactly. For real higher-order signals, we use 339 ocean-drifter trajectories near Madagascar following Schaub et al. (2020), represented as signed edge flows on a complex with 133 vertices, 320 edges, 186 triangles, and $\beta_1 = 2$. The binary clockwise/counter-clockwise labels are determined from trajectory geometry, independently of the harmonic basis used by HARP. We also evaluate partial edge-flow completion on the Anaheim, Barcelona, and Winnipeg transportation networks from the Transportation Networks for Research collection[1] and on the FX300 foreign-exchange network following Jia et al. (2019).

Evaluation protocol. Unless otherwise specified, classification results are reported as mean and standard deviation over stratified 10-fold cross-validation, with identical folds and hyperparameter budgets shared across compared methods. For paired synthetic experiments, both signals in each pair are assigned to the same fold to prevent leakage. For flow completion, all methods use the

[1] https://github.com/bstabler/TransportationNetworks

Table 1: Harmonic information under rank-shared reduction. A/S/T denote the annulus, two-hole strip, and tetrahedron boundary, and $\Delta_{\text{pair}}^{\max}$ measures the maximum distance between harmonic-paired representations.

Representation	Acc. (A / S / T)	$\Delta_{\text{pair}}^{\max}$
Full higher-order	1.0 / 1.0 / 1.0	2.0
Boundary response	0.5/0.5/0.5	$\leq 2.1 \times 10^{-15}$
Schur-reduced vertex	0.5/0.5/0.5	$\leq 2.2 \times 10^{-15}$
Full w/o harmonic	0.5 / 0.5 / 0.5	$\leq 1.2 \times 10^{-15}$
HARP	**1.0 / 1.0 / 1.0**	**2.0**

Table 2: Ocean-drifter classification using different signal representations.

Representation	Dim.	Accuracy
Raw vertex	133	0.811 ± 0.074
Full higher-order	320	0.891 ± 0.046
Schur-reduced	133	0.793 ± 0.062
Raw vertex + HARP	135	0.837 ± 0.062
Reduced + HARP	135	0.846 ± 0.062
Random subspace	2	0.510 ± 0.072
Train-only POD/PCA	2	0.605 ± 0.101
Non-harmonic Hodge	2	0.648 ± 0.115
HARP	2	$\mathbf{0.944 \pm 0.046}$

same observation masks and validation-based regularization search. We compare HARP with the rank-shared reduced representation, the raw higher-order signal, and matched low-dimensional controls, including random, data-driven, and non-harmonic Hodge subspaces. Additional experiments, implementation details, and full results are provided in the Appendix C.

4.2 Structural Validation of Harmonic Information Loss

Synthetic paired signals test whether harmonic information disappears at the rank-shared reduction interface as predicted by Sec. 3. For each synthetic complex, we construct paired signals

$$z_+ = g + c + h, \qquad z_- = g + c - h, \tag{18}$$

where the gradient component g and curl component c are identical within each pair and only the harmonic component h changes sign. We compare the full higher-order signal, its boundary response $B_k z$, the Schur-reduced vertex representation obtained from the reduced system in equation 4, the full signal with its harmonic component removed, and the HARP readout.

As shown in Table 1, the full higher-order signal perfectly distinguishes the paired inputs on all three complexes. Both the boundary response and Schur-reduced vertex representation collapse each harmonic pair to numerical precision, yielding 50% accuracy, as does explicitly removing the harmonic component. In contrast, HARP preserves the full pairwise difference and achieves perfect prediction using only the harmonic coordinates. These results directly verify that the harmonic information lost at the boundary interface remains absent after the full Schur reduction.

4.3 Real-World Information Loss and HARP Recovery

The ocean-drifter task tests whether the same information loss arises in real higher-order signals.

Ocean-drifter classification. Let $z \in \mathbb{R}^{320}$ denote the signed edge flow of a trajectory, and let $T_{\epsilon,1}$ denote the rank-1 block of the higher-order transfer map $T_\epsilon = Q(R + \epsilon I)^{-1}$. Following the reduced system in Sec. 3, we define the reduced vertex signal as

$$x_{\text{red}} = -\left(M_\epsilon^{\downarrow}\right)^{\dagger} T_{\epsilon,1} z. \tag{19}$$

For $z = g + c + h$, both $T_{\epsilon,1} c$ and $T_{\epsilon,1} h$ vanish to numerical precision, and $x_{\text{red}}(z)$ matches $x_{\text{red}}(g)$ accordingly (detailed in Appendix C.3). Thus, the actual rank-shared reduction follows the theoretical prediction: curl and harmonic components do not reach the reduced vertex signal.

Table 2 shows that the harmonic channel substantially improves the reduced representation. Adding only two HARP coordinates increases accuracy from 0.793 to 0.846, an absolute gain of 0.053. Notably, the two-dimensional HARP readout achieves 0.944 accuracy, indicating that task-relevant information in these trajectories is strongly concentrated in the harmonic subspace. This does not imply that HARP contains more information than the full edge signal; rather, it provides a compact representation in which the relevant information is more statistically accessible under finite samples.

Harmonic coordinate interpretation. Figure 2 visualizes the two harmonic directions and the corresponding HARP coordinates on the ocean-drifter complex. The coordinates capture circulation

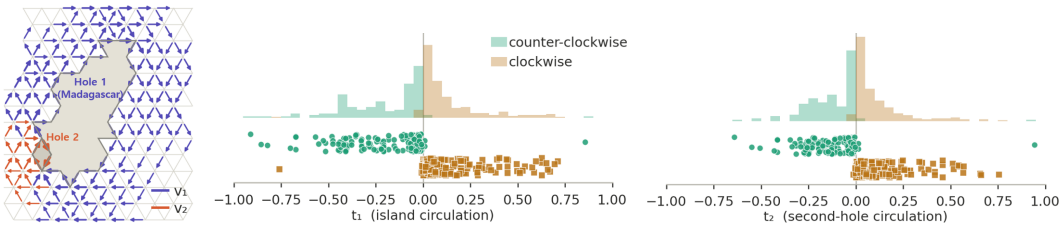

(a) (b) (c)

Figure 2: Harmonic-coordinate interpretation of ocean-drifter trajectories. (a) Harmonic basis of the Madagascar complex with $\beta_1 = 2$, and (b,c) the corresponding HARP coordinates $t_1 = \mathbf{v}_1^\top \mathbf{z}$ and $t_2 = \mathbf{v}_2^\top \mathbf{z}$. For visualization only, the basis is orthogonally rotated within $\mathcal{H}_1$ so that $\mathbf{v}_1$ and $\mathbf{v}_2$ align with circulation around Holes 1 and 2, respectively; this rotation is used only for visualization and does not affect the harmonic subspace or the classification results reported in Table 2. The two coordinates reveal the circulation structure associated with the independent loops of the complex.

Table 3: Partial edge-flow completion across observation ratios. Held-out Pearson correlations are reported; joint results are in Appendix C.5.

Dataset	Method	10%	25%	50%
Anaheim	Gradient	$-0.275_{\pm 0.05}$	$-0.375_{\pm 0.05}$	$-0.385_{\pm 0.07}$
	HARP	$\mathbf{0.463}_{\pm \mathbf{0.05}}$	$\mathbf{0.746}_{\pm \mathbf{0.04}}$	$\mathbf{0.928}_{\pm \mathbf{0.02}}$
Barcelona	Gradient	$-0.242_{\pm 0.02}$	$-0.273_{\pm 0.04}$	$-0.202_{\pm 0.06}$
	HARP	$\mathbf{0.281}_{\pm \mathbf{0.03}}$	$\mathbf{0.453}_{\pm \mathbf{0.03}}$	$\mathbf{0.621}_{\pm \mathbf{0.04}}$
Winnipeg	Gradient	$-0.268_{\pm 0.03}$	$-0.340_{\pm 0.03}$	$-0.342_{\pm 0.04}$
	HARP	$\mathbf{0.183}_{\pm \mathbf{0.05}}$	$\mathbf{0.286}_{\pm \mathbf{0.08}}$	$\mathbf{0.428}_{\pm \mathbf{0.05}}$
FX300	Gradient	$\mathbf{0.947}_{\pm \mathbf{0.08}}$	$\mathbf{1.000}_{\pm \mathbf{0.00}}$	$\mathbf{1.000}_{\pm \mathbf{0.00}}$
	HARP	$-0.760_{\pm 0.07}$	$-0.808_{\pm 0.06}$	$-0.871_{\pm 0.02}$

Table 4: HARP augmentation across lower-rank backbones on ocean drifters.

	Base	+HARP		Base	+HARP
GCN	$0.657_{\pm .087}$	$\mathbf{0.980}_{\pm \mathbf{.023}}$	SAGE	$0.732_{\pm .079}$	$\mathbf{0.979}_{\pm \mathbf{.019}}$
GAT	$0.625_{\pm .108}$	$\mathbf{0.971}_{\pm \mathbf{.037}}$	Cheb	$0.740_{\pm .086}$	$\mathbf{0.974}_{\pm \mathbf{.021}}$
GIN	$0.714_{\pm .076}$	$\mathbf{0.959}_{\pm \mathbf{.041}}$	CEO	$0.723_{\pm .070}$	$\mathbf{0.959}_{\pm \mathbf{.041}}$

Table 5: Effect of harmonic projection in the certified HARP construction on ocean drifters.

Construction	Accuracy
Direct Hodge	0.956 ± 0.039
Morse lift w/o projection	0.546 ± 0.071
HARP (DMT + projection)	0.944 ± 0.046

around the two independent loops and separate clockwise from counter-clockwise trajectories, providing an interpretable view of the harmonic information retained by HARP.

Harmonic specificity. Matched low-dimensional controls confirm that the gain is specific to the harmonic subspace rather than dimensionality alone. The two-dimensional HARP readout achieves 0.944 accuracy, compared with 0.510 for random subspaces, 0.605 for train-only POD/PCA directions, and 0.648 for non-harmonic low-frequency Hodge modes of the same dimension.

4.4 Partial Edge-Flow Completion

We next evaluate HARP when the complete higher-order signal is unavailable at inference time. Given only a subset of observed edges, HARP estimates the harmonic coordinates from the observed flow values and reconstructs the missing edges through the certified harmonic basis. For these datasets, we operate on their 1-skeletons, so the cycle space $\ker B_1$ coincides with the rank-1 harmonic space. This setting is related to divergence-constrained edge-flow recovery (Jia et al., 2019), while here it serves to test whether the harmonic channel identified by our reduction analysis remains informative under partial observation. We compare this reconstruction with a gradient model based on $B_1^\top$; the joint gradient–harmonic objective is reported in Appendix C.5. All methods use identical observation masks, and neither coefficient estimation nor regularization selection accesses the held-out edges.

As shown in Table 3, HARP consistently outperforms the gradient reconstruction on Anaheim, Barcelona, and Winnipeg across all three observation ratios. Its performance improves as more edges are observed, reaching correlations of 0.928, 0.621, and 0.428, respectively, at 50% observation. FX300 exhibits the complementary regime: the gradient reconstruction is already highly predictive, while the harmonic-only reconstruction is strongly anti-correlated with the held-out flows. We use this experiment as a regime diagnostic: harmonic-only reconstruction is informative when the observed flow contains substantial cycle-space variation, whereas gradient-dominated signals such as FX300 exhibit the complementary behavior.

4.5 Backbone and Construction Validation

Backbone augmentation. We examine whether the benefit of exposing harmonic coordinates depends on a lower-rank architecture. As shown in Table 4, the tuned backbones achieve base accuracies between 0.625 and 0.740. Adding HARP consistently improves performance to 0.959–0.980 across spatial, attention-based, spectral, and collapsed-operator architectures, with paired gains of 0.233–0.345 across the tested backbones. This consistent improvement indicates that the harmonic channel provides complementary information beyond the particular lower-rank architecture.

Certified harmonic construction. Table 5 verifies the role of harmonic projection under the same nested 10-fold protocol used in Table 2. The HARP construction achieves 0.944 ± 0.046 accuracy, close to the 0.956 ± 0.039 obtained by direct Hodge decomposition, while the unprojected Morse lift drops to 0.546 ± 0.071. Across six complexes and harmonic ranks 1–3, the projected DMT basis matches the direct Hodge subspace with projector error at most 3.3×10^{-14} and principal cosine 1.0. On grid complexes up to 64^2, DMT reduces the complex to approximately 34–40% of its original size. Sparse Lanczos is substantially faster when only a few null-space vectors are required, while DMT additionally provides a topology-preserving Morse complex and a chain-compatible lift to the original chain space. We therefore use DMT for this structural construction rather than as a faster replacement for specialized sparse eigensolvers. Additional construction, projection, scaling, and basis-ambiguity analyses are reported in Appendix C.

5 Related Works

Higher-Order Signal Processing and Topological Learning. Hodge theory decomposes signals on higher-order structures into gradient, curl, and harmonic components (Lim, 2020; Barbarossa & Sardellitti, 2020). Schaub et al. (2020) further use a normalized Hodge Laplacian to construct edge-space embeddings and show that harmonic projections capture global circulation patterns in ocean-drifter trajectories. Building on these foundations, topological deep learning extends representation learning from graphs to simplicial and cellular complexes through higher-order message passing and attention-based architectures (Bodnar et al., 2021a;b; Hajij et al., 2022; Papamarkou et al., 2024). These approaches analyze or process signals while retaining their higher-order representation. In contrast, HARP asks which signal components become inaccessible after reduction to a lower rank and uses harmonic coordinates specifically to expose the subspace hidden by that reduction.

Higher-Order Reduction and Effective Operators. Schur complements provide a principled mechanism for eliminating variables and constructing effective operators on retained degrees of freedom, with classical applications including Kron reduction of graph Laplacians (Dorfler & Bullo, 2012; Devriendt, 2022). More recently, Collapsed Effective Operators (CEO) apply this principle to higher-order structures by collapsing a graded operator to the vertex level (Krahn et al., 2026). HARP instead characterizes the higher-order information annihilated by such boundary-mediated reductions and explicitly exposes the missing harmonic coordinates.

Topology-Preserving Reduction and Harmonic Representations. Discrete Morse theory provides topology-preserving reductions of cell complexes through chain-homotopy equivalence (Forman, 1998), and has also been studied for signal compression and reconstruction on chain complexes (Ebli et al., 2024). HARP uses Morse reduction to obtain a topology-preserving reduced complex and a chain-compatible lifting map, followed by harmonic projection in the original chain space.

6 Conclusion

We studied what higher-order information remains observable after boundary-mediated reduction and showed that harmonic components are structurally hidden at the observation interface of rank-shared operators. To recover this missing information, we introduced HARP, which combines topology-preserving Morse reduction, chain-compatible lifting, and harmonic projection to construct compact signal-dependent harmonic coordinates. Experiments on synthetic and real higher-order signals validate the predicted information pathway and show that HARP complements lower-rank representations when harmonic information is task-relevant. These results highlight the importance of characterizing not only how higher-order information is aggregated, but also which components remain observable after reduction.

AI USE STATEMENT

We used generative AI tools for language editing, grammar correction, and refinement of manuscript wording during the writing process. Generative AI was not used to generate experimental results, perform data analysis, or determine the scientific conclusions of this work. All AI-assisted text was reviewed and revised by the authors, who take full responsibility for the final content of the paper.

ETHICS STATEMENT

This work focuses on the theoretical and computational analysis of higher-order signal representations and uses existing benchmark datasets. It does not involve human subjects, the collection of personally identifiable or sensitive information, or deployment in safety-critical settings. We are not aware of additional ethical concerns specific to the methods or experiments presented in this work.

REPRODUCIBILITY STATEMENT

We provide the assumptions and complete proofs of the theoretical results in Appendix A, the detailed HARP construction and algorithm in Appendix B, and experimental protocols, implementation details, and additional results in Appendix C. The main paper reports the datasets, evaluation settings, and primary experimental configurations used in our experiments. Code for reproducing the proposed method and experiments will be released upon acceptance.

REFERENCES

Sergio Barbarossa and Stefania Sardellitti. Topological signal processing over simplicial complexes. *IEEE Transactions on Signal Processing*, 2020.

Cristian Bodnar et al. Weisfeiler and lehman go cellular: Cw networks. In *Advances in Neural Information Processing Systems*, 2021a.

Cristian Bodnar et al. Weisfeiler and lehman go topological: Message passing simplicial networks. In *International Conference on Machine Learning*, 2021b.

Karel Devriendt. Effective resistance is more than distance: Laplacians, simplices and the schur complement. *Linear Algebra and its Applications*, 2022.

Florian Dorfler and Francesco Bullo. Kron reduction of graphs with applications to electrical networks. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 2012.

Stefania Ebli et al. Morse theoretic signal compression and reconstruction on chain complexes. *Journal of Applied and Computational Topology*, 2024.

Robin Forman. Morse theory for cell complexes. *Advances in mathematics*, 1998.

Mustafa Hajij et al. Topological deep learning: Going beyond graph data. *arXiv preprint arXiv:2206.00606*, 2022.

Allen Hatcher. *Algebraic topology*. Cambridge university press, 2002.

Junteng Jia et al. Graph-based semi-supervised & active learning for edge flows. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2019.

Maximilian Krahn et al. Collapsed effective operators for higher-order structures. In *International Conference on Machine Learning*, 2026.

Lek-Heng Lim. Hodge laplacians on graphs. *SIAM Review*, 2020.

Eliakim H Moore. On the reciprocal of the general algebraic matrix. *Bulletin of the American Mathematical Society*, 1920.

Theodore Papamarkou et al. Position: Topological deep learning is the new frontier for relational learning. *Proceedings of Machine Learning Research*, 2024.

Mathilde Papillon et al. Architectures of topological deep learning: A survey of message-passing topological neural networks. *arXiv preprint arXiv:2304.10031*, 2023.

Roger Penrose. A generalized inverse for matrices. In *Mathematical Proceedings of the Cambridge Philosophical Society*, 1955.

Michael T Schaub et al. Random walks on simplicial complexes and the normalized hodge 1-laplacian. *SIAM Review*, 2020.

Andrei N Tikhonov. Solution of incorrectly formulated problems and the regularization method. *Sov Dok*, 1963.

A Proofs and Scope of Harmonic Unobservability

This appendix provides the proofs and scope clarifications for the results in Sec. 3.2. We first establish the gradient factorization induced by rank-shared boundary coupling. We then formalize the boundary-compatible coupling class, characterize the effect of cell-dependent anisotropic weights, and prove the kernel-reduction result. Throughout this appendix, the main impossibility result refers specifically to the rank-shared coupling class used in the main text.

A.1 General Setting and Notation

Consider a finite chain complex

$$C_0 \xleftarrow{B_1} C_1 \xleftarrow{B_2} \cdots \xleftarrow{B_K} C_K, \tag{20}$$

with boundary operators satisfying $B_{k-1}B_k = 0$. For each rank k, the Hodge decomposition is

$$C_k = \mathcal{G}_k \oplus \mathcal{C}_k \oplus \mathcal{H}_k, \tag{21}$$

where

$$\mathcal{G}_k = \operatorname{im} B_k^\top, \qquad \mathcal{C}_k = \operatorname{im} B_{k+1}, \qquad \mathcal{H}_k = \ker B_k \cap \ker B_{k+1}^\top. \tag{22}$$

In particular,

$$\ker B_k = \mathcal{C}_k \oplus \mathcal{H}_k. \tag{23}$$

For rank-0 reduction, define the eliminated higher-rank space as

$$\mathcal{E} = \bigoplus_{k=1}^{K} C_k. \tag{24}$$

The boundary-mediated operator in Sec. 3 can then be partitioned as

$$M = \begin{pmatrix} \Delta_0 & Q \\ Q^\top & R \end{pmatrix}, \qquad Q = \begin{bmatrix} -\gamma_1 B_1 & 0 & \cdots & 0 \end{bmatrix}. \tag{25}$$

For $\epsilon > 0$, let

$$R_\epsilon = R + \epsilon I. \tag{26}$$

The corresponding reduced operator and higher-rank source-transfer map are

$$M_\epsilon^{\downarrow} = \Delta_0 - Q R_\epsilon^{-1} Q^\top, \qquad T_\epsilon = Q R_\epsilon^{-1}. \tag{27}$$

A.2 Proof of Theorem 3.1: Gradient Factorization

Let

$$G_\epsilon = \left[R_\epsilon^{-1}\right]_{11} : C_1 \to C_1 \tag{28}$$

denote the rank-1-to-rank-1 diagonal block of R_ϵ^{-1}. Since the only nonzero block of Q acts on C_1,

$$Q R_\epsilon^{-1} Q^\top = \gamma_1^2 B_1 G_\epsilon B_1^\top. \tag{29}$$

Therefore,

$$M_\epsilon^{\downarrow} = \Delta_0 - \gamma_1^2 B_1 G_\epsilon B_1^\top, \tag{30}$$

which proves the first statement of Theorem 3.1. Let $\Pi_{\mathcal{G}_1}$ denote the orthogonal projector onto $\mathcal{G}_1 = \operatorname{im} B_1^\top$. Since $\mathcal{G}_1^\perp = \ker B_1$,

$$B_1 = B_1 \Pi_{\mathcal{G}_1}, \qquad B_1^\top = \Pi_{\mathcal{G}_1} B_1^\top. \tag{31}$$

Hence,

$$M_\epsilon^{\downarrow} = \Delta_0 - \gamma_1^2 B_1 \left(\Pi_{\mathcal{G}_1} G_\epsilon \Pi_{\mathcal{G}_1}\right) B_1^\top. \tag{32}$$

Thus, the Schur correction accesses rank-1 directions only through the reference gradient subspace. We next consider the source-transfer map. The rank-1 space decomposes orthogonally as

$$C_1 = \mathcal{G}_1 \oplus \ker B_1. \tag{33}$$

The within-rank operator Δ_1 preserves this decomposition, while the rank-1-to-rank-2 interaction vanishes on $\mathcal{G}_1$ because

$$B_2^\top B_1^\top = (B_1 B_2)^\top = 0. \tag{34}$$

Hence, $\mathcal{G}_1 \oplus 0 \oplus \cdots \oplus 0$ is invariant under R_ϵ. Since R_ϵ is symmetric, its orthogonal complement

$$\ker B_1 \oplus C_2 \oplus \cdots \oplus C_K \tag{35}$$

is also invariant under R_ϵ, and therefore under R_ϵ^{-1}. Consider a perturbation supported only at rank 1:

$$\delta z = \begin{bmatrix} \delta z_1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \qquad \delta z_1 \in \ker B_1. \tag{36}$$

The rank-1 component of $R_\epsilon^{-1}\delta z$ remains in $\ker B_1$, so

$$T_\epsilon \delta z = Q R_\epsilon^{-1} \delta z = 0. \tag{37}$$

Therefore, if $B_1 z_1 = B_1 z_1'$, then $z_1 - z_1' \in \ker B_1$, and the two sources induce the same retained-level response. Hence, the rank-1 contribution to $T_\epsilon z$ depends on z_1 only through $B_1 z_1$. Since

$$\ker B_1 = \mathcal{C}_1 \oplus \mathcal{H}_1, \tag{38}$$

both curl and harmonic perturbations are unobservable at the rank-0/rank-1 interface. □

A.3 Proof of Theorem 3.2: Boundary-Compatible Coupling Rigidity

We formalize the boundary-compatible coupling conditions used in Theorem 3.2. Let

$$B_1 = [b_1 \quad \cdots \quad b_{n_1}] \tag{39}$$

be the signed incidence matrix, where b_e is the incidence vector of oriented edge e. Consider a linear coupling

$$Y : C_1 \to C_0. \tag{40}$$

Under signed-incidence locality, each edge contribution follows the signed endpoint pattern encoded by b_e:

$$Y(:, e) \in \operatorname{span}\{b_e\}. \tag{41}$$

Hence, there exists a scalar γ_e, possibly determined by intrinsic attributes of edge e, such that

$$Y(:, e) = \gamma_e b_e. \tag{42}$$

Collecting the edge-wise coefficients gives

$$\Gamma = \operatorname{diag}(\gamma_1, \ldots, \gamma_{n_1}), \qquad Y = B_1 \Gamma. \tag{43}$$

Let D_E be a diagonal matrix with entries in $\{-1, +1\}$ representing edge-orientation reversals. Then

$$B_1 \mapsto B_1 D_E, \qquad Y \mapsto Y D_E. \tag{44}$$

Since Γ is diagonal,

$$(B_1 D_E)\Gamma = B_1 \Gamma D_E = Y D_E, \tag{45}$$

so the edge-wise weights are independent of the arbitrary orientation convention. Likewise, for vertex and edge permutation matrices P_V and P_E,

$$B_1 \mapsto P_V B_1 P_E^\top, \qquad Y \mapsto P_V Y P_E^\top. \tag{46}$$

Relabeling equivariance therefore attaches each γ_e to the corresponding edge and its intrinsic attributes rather than to an arbitrary matrix index. If the coupling rule is independent of cell-specific attributes and shares one coupling strength across rank-1 cells, then

$$\gamma_e = \gamma \qquad \text{for all } e. \tag{47}$$

Consequently,

$$\Gamma = \gamma I, \qquad Y = \gamma B_1. \tag{48}$$

This recovers the rank-shared boundary-mediated coupling class used in the main analysis. □

A.4 Anisotropic Coupling and Covariant Unobservability

Consider the boundary-compatible coupling

$$Y = B_1\Gamma, \tag{49}$$

where Γ is a diagonal matrix of edge-wise weights. Assume first that Γ is invertible. Then

$$x \in \ker(B_1\Gamma) \iff \Gamma x \in \ker B_1, \tag{50}$$

and therefore

$$\ker(B_1\Gamma) = \Gamma^{-1}\ker(B_1). \tag{51}$$

Since multiplication by an invertible Γ is bijective,

$$\dim\ker(B_1\Gamma) = \dim\ker B_1. \tag{52}$$

Thus, anisotropic reweighting preserves the dimension of the blind subspace but generally changes its orientation. For the rank-shared case $\Gamma = \gamma I$ with $\gamma \neq 0$,

$$\ker(\gamma B_1) = \ker B_1. \tag{53}$$

If Γ is singular, its own kernel can introduce additional blind directions and increase the nullity. Let $\Pi_{\mathcal{G}_1}$ denote the orthogonal projector onto the reference gradient subspace. Since

$$B_1 = B_1\Pi_{\mathcal{G}_1}, \tag{54}$$

every $h \in \mathcal{H}_1$ satisfies

$$B_1\Gamma h = B_1\Pi_{\mathcal{G}_1}\Gamma h. \tag{55}$$

Hence, a reference-harmonic direction can become observable when Γh acquires a nonzero component in the reference gradient subspace. An anisotropic coupling therefore re-encodes reference-harmonic variation into a boundary-visible direction, whereas HARP directly reads the harmonic coordinates defined by the original reference Hodge geometry.

Covariant unobservability. Define

$$\widetilde{B}_1 = B_1\Gamma, \qquad \widetilde{B}_2 = \Gamma^{-1}B_2. \tag{56}$$

Theorem A.1 (Covariant Unobservability)**.** *Let Γ be an invertible diagonal matrix. Then $(\widetilde{B}_1, \widetilde{B}_2)$ defines a chain complex that is chain-isomorphic to the reference complex and therefore preserves the first Betti number β_1. Let*

$$\widetilde{\mathcal{H}}_1 = \ker\widetilde{B}_1 \cap \ker\widetilde{B}_2^{\top}. \tag{57}$$

Then

$$\dim\widetilde{\mathcal{H}}_1 = \beta_1, \qquad \widetilde{B}_1\widetilde{h} = 0 \quad \textit{for every } \widetilde{h} \in \widetilde{\mathcal{H}}_1. \tag{58}$$

Hence, anisotropic reweighting may expose directions that are harmonic under the reference geometry while remaining blind to the harmonic subspace defined by its own reweighted geometry.

Proof. The transformed boundary operators satisfy

$$\widetilde{B}_1\widetilde{B}_2 = B_1\Gamma\Gamma^{-1}B_2 = B_1B_2 = 0. \tag{59}$$

Thus, they define a valid chain complex. Because Γ is invertible, the transformation is an invertible change of coordinates on C_1, so the transformed and reference complexes are chain-isomorphic and have identical homology dimensions. Therefore,

$$\dim\widetilde{\mathcal{H}}_1 = \beta_1. \tag{60}$$

Finally, every $\widetilde{h} \in \widetilde{\mathcal{H}}_1$ belongs to $\ker\widetilde{B}_1$, and hence

$$\widetilde{B}_1\widetilde{h} = 0. \tag{61}$$

□

A.5 Proof of Theorem 3.3: Kernel Reduction

Let

$$\mathcal{K}_H = \bigoplus_{k=1}^{K} \mathcal{H}_k \tag{62}$$

denote the direct sum of the harmonic subspaces contained in the eliminated ranks. For every $h_k \in \mathcal{H}_k$,

$$B_k h_k = 0, \qquad B_{k+1}^\top h_k = 0, \qquad \Delta_k h_k = 0. \tag{63}$$

The within-rank block and both adjacent boundary-mediated interactions therefore vanish on h_k. Hence,

$$\mathcal{K}_H \subseteq \ker R. \tag{64}$$

Let

$$U_H = [u_1 \quad \cdots \quad u_\beta], \qquad \beta = \sum_{k=1}^{K} \beta_k, \tag{65}$$

be an orthonormal basis of $\mathcal{K}_H$. Since the rank-1 component of every vector in $\mathcal{K}_H$ belongs to $\mathcal{H}_1 \subseteq \ker B_1$,

$$QU_H = 0. \tag{66}$$

For the Moore–Penrose pseudoinverse,

$$R^\dagger U_H = 0. \tag{67}$$

For Tikhonov regularization,

$$(R + \epsilon I)^{-1} U_H = \epsilon^{-1} U_H, \tag{68}$$

but

$$Q(R + \epsilon I)^{-1} U_H = \epsilon^{-1} Q U_H = 0. \tag{69}$$

Finally, consider assigning finite positive mass to the harmonic directions:

$$R_A = R + U_H A U_H^\top, \qquad A \succ 0. \tag{70}$$

On $\mathcal{K}_H$, this gives the finite inverse response A^{-1}, but its contribution to the retained-space correction is

$$QU_H A^{-1} U_H^\top Q^\top = 0. \tag{71}$$

Thus, the Moore–Penrose pseudoinverse, Tikhonov regularization, and finite positive harmonic masses differ internally but yield the same retained-space outcome. In every case, the harmonic contribution vanishes because $QU_H = 0$. Hence, harmonic information is lost at the observation interface rather than through the numerical treatment of the singular higher-rank block. □

A.6 Scope Beyond the Rank-Shared Class

The preceding results do not imply that harmonic information is inaccessible to every possible linear or nonlinear measurement. For example,

$$Y_{\mathcal{H}} = V_{\mathcal{H},1}^\top \tag{72}$$

directly returns the harmonic coordinates of a rank-1 signal. Such a global measurement is not a local boundary-mediated coupling and therefore lies outside the operator class considered in Theorems 3.1–3.3. Likewise, Appendix A.4 shows that a cell-dependent coupling $B_1\Gamma$ can expose variation that is harmonic under the reference Hodge geometry by re-encoding it into boundary-visible directions. The scope of the main impossibility result is therefore precise: rank-shared boundary-mediated coupling preserves the reference boundary null space, whereas more general measurement rules can alter which higher-order directions are observable. HARP addresses this blind spot by explicitly reading the reference-harmonic coordinates without modifying the underlying boundary-mediated observation interface.

B TOPOLOGY-PRESERVING MORSE REDUCTION AND HARMONIC CONSTRUCTION

This appendix details the constructive part of HARP. The DMT pipeline assumes a finite cellular realization for which a Forman acyclic matching and the corresponding Morse complex are defined. We first describe the discrete Morse reduction, then establish the chain-compatible lifting property, and finally show why harmonic projection recovers the reference-harmonic subspace of the original complex.

B.1 DISCRETE MORSE REDUCTION

Let

$$(C_*, B_*) \tag{73}$$

denote the original finite chain complex. Forman discrete Morse theory constructs an acyclic matching between incident cells of adjacent ranks, while unmatched cells are retained as critical cells. The critical rank-k cells span the Morse chain space C_k^M, yielding

$$C_0^M \xleftarrow{B_1^M} C_1^M \xleftarrow{B_2^M} \cdots \xleftarrow{B_K^M} C_K^M. \tag{74}$$

The Morse complex is chain-homotopy equivalent to the original complex and therefore preserves homology:

$$H_k(C_*^M) \cong H_k(C_*), \qquad \beta_k^M = \beta_k. \tag{75}$$

HARP keeps all rank-0 cells critical, so

$$C_0^M = C_0. \tag{76}$$

This preserves the original vertex domain while allowing higher-rank cells to be reduced through the Morse matching.

B.2 CHAIN-COMPATIBLE LIFTING

A critical-cell coordinate is defined in the Morse complex and should not be identified directly with a coordinate in the original chain space. HARP uses the lifting map

$$E_k : C_k^M \to C_k, \tag{77}$$

which satisfies

$$B_k E_k = E_{k-1} B_k^M. \tag{78}$$

Let $z^M \in C_k^M$ be a cycle, so $B_k^M z^M = 0$. Then

$$B_k E_k z^M = E_{k-1} B_k^M z^M = 0. \tag{79}$$

Hence,

$$E_k z^M \in \ker B_k. \tag{80}$$

Thus, the lifting map preserves the cycle property required for subsequent harmonic projection.

B.3 FROM MORSE HARMONICS TO REFERENCE HARMONICS

For each rank k, the Morse Hodge Laplacian is

$$\Delta_k^M = (B_k^M)^\top B_k^M + B_{k+1}^M (B_{k+1}^M)^\top. \tag{81}$$

Let $V_{\mathcal{H},k}^M$ be an orthonormal basis of $\ker \Delta_k^M$. Since the Morse and original complexes are chain-homotopy equivalent,

$$\dim \ker \Delta_k^M = \beta_k. \tag{82}$$

We lift the Morse harmonic basis as

$$V_{\mathrm{lift},k} = E_k V_{\mathcal{H},k}^M. \tag{83}$$

By equation 78,

$$B_k V_{\mathrm{lift},k} = 0. \tag{84}$$

However, E_k preserves homology rather than the Euclidean inner product of the original chain space, so a lifted cycle can contain a component in $\operatorname{im} B_{k+1}$. The orthogonal projector onto this boundary subspace is

$$P_{\mathcal{C}_k} = B_{k+1} B_{k+1}^{\dagger}. \tag{85}$$

We therefore define

$$\widetilde{V}_{\mathcal{H},k} = \left(I - B_{k+1} B_{k+1}^{\dagger} \right) E_k V_{\mathcal{H},k}^{M}. \tag{86}$$

Since $B_k V_{\mathrm{lift},k} = 0$ and $B_k B_{k+1} = 0$,

$$B_k \widetilde{V}_{\mathcal{H},k} = 0. \tag{87}$$

Moreover,

$$B_{k+1}^{\top} \widetilde{V}_{\mathcal{H},k} = 0. \tag{88}$$

Hence,

$$\operatorname{span}(\widetilde{V}_{\mathcal{H},k}) \subseteq \mathcal{H}_k. \tag{89}$$

It remains to show that the projection does not remove a nontrivial homology class. Suppose that

$$\widetilde{V}_{\mathcal{H},k} a = 0 \tag{90}$$

for some coefficient vector a. Then

$$E_k V_{\mathcal{H},k}^{M} a \in \operatorname{im} B_{k+1}, \tag{91}$$

so the corresponding lifted cycle represents the zero homology class in the original complex. Because E_k realizes the chain-homotopy equivalence, its induced map on homology is an isomorphism. Hence, $V_{\mathcal{H},k}^{M} a$ also represents the zero homology class in the Morse complex. Since the columns of $V_{\mathcal{H},k}^{M}$ represent a basis of the β_k-dimensional Morse homology, this implies $a = 0$. Therefore,

$$\operatorname{rank}(\widetilde{V}_{\mathcal{H},k}) = \beta_k. \tag{92}$$

Together with $\dim \mathcal{H}_k = \beta_k$,

$$\operatorname{span}(\widetilde{V}_{\mathcal{H},k}) = \mathcal{H}_k. \tag{93}$$

Finally,

$$V_{\mathcal{H},k} = \operatorname{orth}(\widetilde{V}_{\mathcal{H},k}), \tag{94}$$

which satisfies

$$B_k V_{\mathcal{H},k} = 0, \qquad B_{k+1}^{\top} V_{\mathcal{H},k} = 0. \tag{95}$$

B.4 Basis Ambiguity and Invariant Harmonic Content

The harmonic subspace $\mathcal{H}_k$ is uniquely defined, but its orthonormal basis is not. For any orthogonal matrix $O \in \mathbb{R}^{\beta_k \times \beta_k}$,

$$V'_{\mathcal{H},k} = V_{\mathcal{H},k} O \tag{96}$$

spans the same harmonic subspace. The corresponding coordinate vector transforms as

$$t' = (V'_{\mathcal{H},k})^{\top} z = O^{\top} t. \tag{97}$$

However, the represented harmonic component is basis invariant:

$$V'_{\mathcal{H},k} t' = V_{\mathcal{H},k} O O^{\top} V_{\mathcal{H},k}^{\top} z = V_{\mathcal{H},k} V_{\mathcal{H},k}^{\top} z. \tag{98}$$

Thus, HARP identifies a unique harmonic projection even though its coordinate representation is defined only up to an orthogonal transformation within $\mathcal{H}_k$. For the single-complex experiments in Sec. 4, we construct one deterministic basis and reuse it for all samples. We additionally canonicalize the sign of each basis vector by requiring its first numerically nonzero entry to be positive. This removes independent sign flips but does not resolve the general $O(\beta_k)$ rotational ambiguity when $\beta_k > 1$. The effect of this ambiguity at larger harmonic dimension is examined in Appendix C.11.

Algorithm 1 HARP harmonic basis construction

Require: Finite chain complex (C_*, B_*), ranks of interest $\mathcal{K}$
Ensure: Harmonic bases $\{V_{\mathcal{H},k}\}_{k\in\mathcal{K}}$
1: Construct an acyclic Forman matching on ranks $k \geq 1$, keeping all rank-0 cells critical.
2: Build the Morse complex (C_*^M, B_*^M) and chain-compatible lifting maps E_k.
3: **for** $k \in \mathcal{K}$ **do**
4: $\quad \Delta_k^M \leftarrow (B_k^M)^\top B_k^M + B_{k+1}^M (B_{k+1}^M)^\top$
5: $\quad$ Compute an orthonormal basis $V_{\mathcal{H},k}^M$ of $\ker \Delta_k^M$.
6: $\quad V_{\mathrm{lift},k} \leftarrow E_k V_{\mathcal{H},k}^M$
7: $\quad \widetilde{V}_{\mathcal{H},k} \leftarrow (I - B_{k+1} B_{k+1}^\dagger) V_{\mathrm{lift},k}$
8: $\quad V_{\mathcal{H},k} \leftarrow \mathrm{orth}(\widetilde{V}_{\mathcal{H},k})$
9: **end for**
10: **return** $\{V_{\mathcal{H},k}\}_{k\in\mathcal{K}}$

B.5 HARP Construction and Readout

Algorithm 1 summarizes the harmonic basis construction described above. In implementation, we verify the chain identities $B_{k-1}^M B_k^M = 0$ and $B_k E_k = E_{k-1} B_k^M$, preservation of the relevant Betti numbers, and the residual conditions $B_k V_{\mathcal{H},k} \approx 0$ and $B_{k+1}^\top V_{\mathcal{H},k} \approx 0$. For deterministic coordinate reporting, we additionally fix the sign of each basis vector as described in Appendix B.4. Once the basis is constructed, obtaining the HARP representation requires no further optimization. Collecting the rank-wise bases as

$$V_{\mathcal{H}} = \mathrm{blkdiag}(V_{\mathcal{H},1}, \ldots, V_{\mathcal{H},K}), \tag{99}$$

the harmonic coordinates are

$$t(z) = V_{\mathcal{H}}^\top z. \tag{100}$$

C Experimental Details and Additional Results

This appendix provides the experimental protocols, auxiliary controls, and mechanism-level analyses supporting Sec. 4.

C.1 Evaluation Protocols

Unless otherwise specified, classification results use stratified 10-fold cross-validation with identical folds and matched hyperparameter budgets across compared methods. For paired synthetic experiments, both members of each harmonic pair are assigned to the same fold to prevent information leakage. For partial edge-flow completion, all methods use identical observation masks, and coefficient estimation and regularization selection use only the observed edges. Held-out edges are reserved exclusively for final evaluation. Several mechanism-specific analyses use dedicated repeated-seed protocols because they address questions distinct from the main classification evaluation. These protocols are stated explicitly in the corresponding subsections, and their variability estimates are not mixed with the main cross-validation statistics.

C.2 Structural Validation Details

For each synthetic complex, we construct paired signals

$$z^+ = g + c + h, \qquad z^- = g + c - h, \tag{101}$$

where $g \in \mathcal{G}_k$, $c \in \mathcal{C}_k$, and $h \in \mathcal{H}_k$. The class-dependent difference is

$$z^+ - z^- = 2h. \tag{102}$$

Train/test partitioning is performed at the pair level. For a representation ϕ, define

$$\Delta_{\text{pair}} = \|\phi(z^+) - \phi(z^-)\|_2, \tag{103}$$

and

$$\Delta_{\text{pair}}^{\max} = \max_{\text{pairs}} \Delta_{\text{pair}}. \tag{104}$$

For the boundary response,

$$B_k z^+ - B_k z^- = 2B_k h = 0. \tag{105}$$

The rank-shared Schur-reduced representation likewise removes the harmonic difference. In contrast,

$$V_{\mathcal{H},k}^\top z^+ - V_{\mathcal{H},k}^\top z^- = 2V_{\mathcal{H},k}^\top h. \tag{106}$$

Thus, the pairwise distances reported in Table 1 directly test the predicted information pathway independently of downstream classifier capacity.

C.3 Ocean-Drifter Reduction Certificates

Let

$$z = g + c + h \tag{107}$$

denote the Hodge decomposition of an ocean-drifter edge flow. For

$$T_\epsilon = Q(R + \epsilon I)^{-1}, \tag{108}$$

let $T_{\epsilon,1}$ denote the block acting on the rank-1 signal. The implemented reduction satisfies

$$\|T_{\epsilon,1} c\|_2 \le 9.7 \times 10^{-16}, \tag{109}$$

and

$$\|T_{\epsilon,1} h\|_2 \le 5.7 \times 10^{-15}. \tag{110}$$

Using

$$x_{\text{red}} = -(M_\epsilon^{\downarrow})^\dagger T_{\epsilon,1} z, \tag{111}$$

we additionally obtain

$$\|x_{\text{red}}(z) - x_{\text{red}}(g)\|_2 \le 1.2 \times 10^{-14}. \tag{112}$$

These certificates verify on the real complex that the implemented reduction follows the information pathway characterized by Theorem 3.1. The reduced signal x_{red} is distinct from the direct boundary response $B_1 z$: the former is obtained after propagation through the full reduced system.

C.4 Matched Low-Dimensional Controls

All matched controls have dimension two, equal to the certified harmonic dimension $\beta_1 = 2$ of the ocean-drifter complex, and use the same outer stratified 10-fold splits as the main experiment. Within each outer fold, ridge regularization is selected using an inner stratified 80/20 split over

$$\lambda \in \{10^{-4}, 10^{-3}, \dots, 10^{4}\}. \tag{113}$$

The random control uses a fixed two-dimensional random subspace. The POD/PCA basis is fit using only the training pool of each outer fold. The non-harmonic Hodge control uses neighboring low-frequency Hodge directions of the same dimension. HARP achieves 0.944 ± 0.046 accuracy, compared with 0.510 ± 0.072 for the random subspace, 0.605 ± 0.101 for train-only POD/PCA, and 0.648 ± 0.115 for the non-harmonic Hodge control. The paired fold-wise improvements are 0.434 ± 0.062, 0.339 ± 0.089, and 0.297 ± 0.144, respectively, with HARP outperforming each control on all 10 outer folds.

C.5 Partial Edge-Flow Completion

Let P_Ω denote the selector for the observed edge set Ω, and let

$$z_\Omega = P_\Omega z. \tag{114}$$

All coefficient fitting and regularization selection use only observed edges. The held-out edges are used exclusively for final evaluation. The gradient baseline estimates

$$\hat{a} = \arg\min_a \|P_\Omega B_1^\top a - z_\Omega\|_2^2 + \lambda_{\mathcal{G}} \|a\|_2^2, \tag{115}$$

with

$$\hat{z}_{\mathcal{G}} = B_1^\top \hat{a}. \tag{116}$$

HARP estimates

$$\hat{t} = \arg\min_t \|P_\Omega V_{\mathcal{H}} t - z_\Omega\|_2^2 + \lambda_{\mathcal{H}} \|t\|_2^2, \tag{117}$$

with

$$\hat{z}_{\mathcal{H}} = V_{\mathcal{H}} \hat{t}. \tag{118}$$

The joint model solves

$$(\hat{a}, \hat{t}) = \arg\min_{a,t} \|P_\Omega (B_1^\top a + V_{\mathcal{H}} t) - z_\Omega\|_2^2 + \lambda_{\mathcal{G}} \|a\|_2^2 + \lambda_{\mathcal{H}} \|t\|_2^2. \tag{119}$$

The regularization coefficients are selected using a validation split drawn only from the observed edges. These experiments use the 1-skeletons of the networks, so

$$\mathcal{H}_1 = \ker B_1. \tag{120}$$

Across the four flow-completion complexes,

$$\|B_1 V_{\mathcal{H}}\|_2 \le 2.5 \times 10^{-14}, \tag{121}$$

and the DMT-constructed subspace has principal-subspace agreement 1.000000 with the direct null-space construction.

C.6 Testing the Rank-Sharing Assumption

The main theorem assumes a single shared coupling strength within each rank. We test the effect of relaxing this assumption by comparing

$$Y = B_1, \qquad Y = \gamma B_1, \qquad Y = B_1 \Gamma \tag{122}$$

on the reference-harmonic parity task. The experiment uses 48 harmonic pairs with noise level 0.05, pair-safe train/test splitting, and 10 independent seeds. The rank-shared model learns one positive scalar γ, whereas the anisotropic model learns one diagonal weight per edge. We define

$$r_{\mathcal{H}} = \frac{\|Yh\|_2}{\|Y\|_2}. \tag{123}$$

Table 6: Partial edge-flow completion across observation ratios. Values are held-out Pearson correlations, reported as mean ± standard deviation over 10 masks/seeds. Gradient+HARP denotes the joint objective in Equation 119.

Dataset	Observed	Gradient	HARP	Gradient+HARP
Anaheim	10%	-0.275 ± 0.051	0.463 ± 0.054	0.323 ± 0.296
	25%	-0.375 ± 0.045	0.746 ± 0.037	0.734 ± 0.063
	50%	-0.385 ± 0.072	0.928 ± 0.016	0.928 ± 0.016
Barcelona	10%	-0.242 ± 0.024	0.281 ± 0.033	0.256 ± 0.048
	25%	-0.273 ± 0.044	0.453 ± 0.033	0.452 ± 0.031
	50%	-0.202 ± 0.056	0.621 ± 0.038	0.617 ± 0.033
Winnipeg	10%	-0.268 ± 0.033	0.183 ± 0.046	0.181 ± 0.046
	25%	-0.340 ± 0.032	0.286 ± 0.079	0.254 ± 0.071
	50%	-0.342 ± 0.041	0.428 ± 0.053	0.411 ± 0.047
FX300	10%	0.947 ± 0.083	-0.760 ± 0.069	0.946 ± 0.084
	25%	1.000 ± 0.000	-0.808 ± 0.057	1.000 ± 0.000
	50%	1.000 ± 0.000	-0.871 ± 0.024	1.000 ± 0.000

Table 7: Boundary of the rank-sharing assumption. Cell-dependent anisotropic weights can expose reference-harmonic variation, whereas learning a shared scalar cannot.

Complex	Coupling	Accuracy	$r_{\mathcal{H}}$
Annulus	Fixed B_1	0.510 ± 0.017	4.5×10^{-16}
Annulus	Learned γB_1	0.500 ± 0.023	4.7×10^{-16}
Annulus	Learned $B_1\Gamma$	$\mathbf{0.983 \pm 0.026}$	0.44
Two-hole strip	Fixed B_1	0.502 ± 0.032	9.2×10^{-16}
Two-hole strip	Learned γB_1	0.504 ± 0.038	9.4×10^{-16}
Two-hole strip	Learned $B_1\Gamma$	$\mathbf{0.990 \pm 0.010}$	0.39

For every learned anisotropic coupling, the projector distance between $\ker(B_1\Gamma)$ and $\Gamma^{-1}\ker B_1$ is at most 4.6×10^{-15}. By contrast, the projector distance between $\ker(B_1\Gamma)$ and the unweighted reference kernel reaches approximately 0.92, while the nullity remains unchanged. We additionally verify

$$B_1\Gamma h = B_1\Pi_{\mathcal{G}_1}\Gamma h \tag{124}$$

with residual below 10^{-15}. These results empirically match Appendix A.4: within this coupling family, rank sharing preserves the reference blind subspace, whereas cell-dependent weighting can transmit reference-harmonic information through boundary-visible directions.

C.7 HARP versus Learned Anisotropic Re-Encoding

An anisotropic boundary operator provides an alternative route for exposing reference-harmonic variation. We compare this learned re-encoding strategy with HARP on ocean drifters. All representations use the same stratified split at each training fraction, a shared linear logistic head, learning rates in $\{10^{-2}, 10^{-3}\}$, Adam with weight decay 5×10^{-4}, and 10 independent seeds. The anisotropic model learns edge-specific coupling weights from task supervision, whereas HARP uses the fixed two-dimensional certified harmonic readout.

At the 5% training fraction, HARP exceeds the learned anisotropic model by 0.214 accuracy. At the 80% fraction, HARP achieves 0.876 accuracy and 0.992 ROC-AUC, compared with 0.818 and 0.901 for the learned anisotropic model. The learned anisotropic model also confirms the predicted mechanism. Its reference-harmonic observability grows from approximately 3.7×10^{-15} for the scalar coupling to 0.042 and 0.131 at the 5% and 80% training fractions. The corresponding reference-gradient mixing increases to 0.038 and 0.141, respectively. Thus, the learned model exposes reference-harmonic variation through a task-specific re-encoding, whereas HARP reads the reference-harmonic coordinates explicitly without learning that transformation.

Table 8: Sample-efficiency comparison between explicit HARP coordinates and learned anisotropic re-encoding on ocean drifters. The parameter count includes the prediction head. ROC-AUC is reported at the 80% training fraction.

Method	Params.	5%	10%	20%	40%	80%	AUC
Learned $B_1\Gamma$	454	0.659 ± 0.039	0.637	0.713	0.761	0.818 ± 0.074	0.901
HARP readout	3	$\mathbf{0.873 \pm 0.066}$	0.876	0.908	0.860	$\mathbf{0.876 \pm 0.045}$	0.992
Raw edge reference	321	0.717	0.733	0.794	0.832	0.874	0.942

Table 9: Reference-harmonic observability O and gradient mixing D under increasing anisotropy. Values are means over five random anisotropy directions. The identity $B_1\Gamma V_{\mathcal{H}} = B_1\Pi_{\mathcal{G}_1}\Gamma V_{\mathcal{H}}$ holds to machine precision.

Complex	t	$\mathrm{cond}(\Gamma)$	$O(\Gamma)$	$D(\Gamma)$
Annulus	0.00	1.0	0.000	0.000
Annulus	0.75	3.4	0.138	0.196
Annulus	1.50	11.9	0.194	0.324
Annulus	2.25	42.2	0.216	0.398
Annulus	3.00	152.4	0.228	0.443
Two-hole strip	0.00	1.0	0.000	0.000
Two-hole strip	0.75	3.8	0.180	0.186
Two-hole strip	1.50	14.2	0.268	0.328
Two-hole strip	2.25	54.5	0.304	0.425
Two-hole strip	3.00	210.4	0.312	0.488
Ocean drifters	0.00	1.0	0.000	0.000
Ocean drifters	0.75	3.9	0.146	0.183
Ocean drifters	1.50	15.6	0.222	0.346
Ocean drifters	2.25	62.2	0.241	0.469
Ocean drifters	3.00	248.7	0.232	0.551

C.8 Observability versus Reference-Hodge Mixing

To isolate this mechanism, we construct

$$\Gamma(t) = \mathrm{diag}(\exp(tu)), \tag{125}$$

where u is a fixed random direction and $t \geq 0$ controls the degree of anisotropy. We repeat the experiment over five random directions. We measure

$$O(\Gamma) = \frac{\|B_1\Gamma V_{\mathcal{H}}\|_F}{\|B_1\Gamma\|_2}, \tag{126}$$

and

$$D(\Gamma) = \frac{\|\Pi_{\mathcal{G}_1}\Gamma V_{\mathcal{H}}\|_F}{\|\Gamma V_{\mathcal{H}}\|_F}. \tag{127}$$

Both quantities vanish at $t = 0$, corresponding to the rank-shared operator. As anisotropy increases, reference-harmonic observability emerges together with gradient mixing. The maximum residual of the exact identity remains below 3.7×10^{-15}. This supports the algebraic interpretation in equation 55.

C.9 Certified DMT Construction and Harmonic Projection

Across six tested complexes and harmonic ranks 1–3, the projected DMT basis agrees with direct Hodge decomposition with projector error at most 3.3×10^{-14} and minimum principal cosine 1.0 to numerical precision. On ocean drifters, the projected basis has normalized cocycle residual 1.25×10^{-14}, whereas the unprojected Morse lift has residual 1.465. These measurements confirm that the Morse lift preserves the relevant homology class but need not itself be harmonic under

Table 10: Construction cost on the largest 64^2 grid. The DMT basis-only path contains Morse reduction, structural certification, Morse Hodge construction, lifting, and projection. The full-build row additionally assembles operators not required to obtain $V_{\mathcal{H}}$.

Construction	Time	Peak memory
Direct sparse Lanczos + Betti computation	0.71 s	1.12 GB
DMT basis-only	17.5 s	1.12 GB
Direct dense eigendecomposition	86.0 s	4.71 GB
DMT full build	196.8 s	9.34 GB

the reference inner product. Under the same nested 10-fold protocol as Table 2, direct Hodge decomposition achieves 0.956 ± 0.039 accuracy, HARP achieves 0.944 ± 0.046, and the unprojected Morse lift achieves 0.546 ± 0.071. The projected HARP and direct-Hodge bases span the same harmonic subspace with minimum principal cosine 1.0, despite differing by an orthogonal rotation within that subspace. The small downstream difference is therefore consistent with coordinate-dependent conditioning under standardization and regularization rather than a difference in the represented harmonic subspace. This result isolates the role of projection: DMT preserves homology, while projection onto the orthogonal complement of $\mathrm{im}\, B_{k+1}$ recovers the reference-harmonic representative.

C.10 Scaling of the Certified Construction

Across the tested grid complexes, the Morse complex retains approximately 34–40% of the original number of cells. Table 10 reports the largest 64^2 grid, containing 24,064 cells.

Sparse Lanczos is substantially faster when only a few null-space vectors are required. DMT instead provides a reduced chain complex, topology-preserving correspondence, and lifting certificate together with the harmonic interface. Compared with direct dense eigendecomposition, the basis-only DMT pipeline is $4.9\times$ faster at the largest scale. The subspace-fidelity guard remains below 7.8×10^{-8} throughout the scaling suite.

C.11 Large-β Basis Ambiguity

The primary real-data experiment has $\beta_1 = 2$, where a single deterministic basis is constructed once and reused across all samples. Sign canonicalization removes independent sign flips, while the general $O(\beta)$ rotational ambiguity remains. We additionally stress-test the readout in a synthetic setting with $\beta = 121$. Using all sign-canonical coordinates yields

$$0.930 \pm 0.053 \tag{128}$$

accuracy. Full-dimensional varimax and train-only PCA coordinates obtain 0.883 and 0.633, respectively. Selecting informative coordinates using training data after canonical rotation improves performance: both varimax-top-32 and train-only-PCA-top-32 reach 1.000 accuracy, while the corresponding raw 1512-dimensional signal reaches 0.963. These results indicate that large harmonic dimension can require an additional coordinate-selection mechanism. They do not establish a universal canonical basis for $\beta > 1$. The general $O(\beta)$ basis ambiguity therefore remains a limitation of coordinate-level harmonic representations, although the projected harmonic signal

$$V_{\mathcal{H}} V_{\mathcal{H}}^{\top} z \tag{129}$$

is basis invariant.

C.12 Backbone Augmentation Details

The backbone experiment in Sec. 4.5 holds the higher-order information channel fixed while varying the lower-rank architecture. Every lower-rank backbone receives the same boundary-visible input $B_1 z$. The HARP-augmented version additionally receives the same sign-canonical harmonic coordinate vector. No backbone is allowed to construct a different task-specific harmonic basis. The harmonic vector is broadcast to the retained vertex domain and concatenated at the model-level interface

described in Sec. 3.4. The reported implementation adds 48 trainable parameters to the backbone-level model. This parameter count is distinct from the three parameters reported for the two-dimensional HARP classifier in Appendix C.7, since the latter is a separate sample-efficiency diagnostic with a two-weight-plus-bias prediction head.

D SCOPE AND LIMITATIONS

Operator scope. Our unobservability result applies to the rank-shared boundary-mediated operator class analyzed in Sec. 3. It does not imply that harmonic information is inaccessible to arbitrary measurements or more general cross-rank couplings. As shown in Appendix A.4, cell-dependent anisotropic weights can expose variation that is harmonic under the reference geometry by re-encoding it into boundary-visible directions. HARP instead preserves the reference measurement geometry and explicitly reads the harmonic coordinates hidden from the rank-shared interface.

Task dependence. HARP is useful when harmonic variation is relevant to the downstream signal or task, but it is not expected to improve representations in every regime. The FX300 flow-completion experiment provides a complementary case in which the gradient channel is already highly predictive while the harmonic-only reconstruction is uninformative. Moreover, when $\beta_k = 0$, the corresponding harmonic subspace is trivial and HARP provides no additional coordinates at rank k.

Basis ambiguity. Although the harmonic subspace and projected harmonic signal are basis invariant, the coordinate vector $t(z)$ is defined only up to an orthogonal transformation within $\mathcal{H}_k$. We use a deterministic basis with sign canonicalization for the experiments in this work, while the general $O(\beta_k)$ ambiguity remains when $\beta_k > 1$. Appendix C.11 examines this limitation empirically and indicates that large harmonic dimension may require additional coordinate selection.

Construction and input scope. Direct Hodge decomposition can recover the same reference-harmonic subspace when only the basis itself is required. The role of DMT in HARP is instead to provide a topology-preserving reduced complex together with a chain-compatible correspondence to the original chain space, rather than to serve as a faster eigensolver. The current formulation also assumes access to the relevant higher-rank signal, either fully or through partial higher-rank observations as in the completion experiments.