LOGOS: RESIDUAL-GATED WITNESS REFINEMENT FOR NON-RIGID SHAPE CORRESPONDENCE

Anonymous authors
Paper under double-blind review

ABSTRACT

A correspondence sends every vertex of one mesh to a vertex of another. Truncated spectral bases blur small parts and still swap left with right. An update that moves each vertex on its own can tear one region into islands. A network trained on one connectivity is accurate there and degrades when the mesh changes. LOGOS is a training-free editor of a hard correspondence. A residual gate reassigns only high-residual vertices, cleanup freezes the largest piece of each cluster, and a polish moves points inside that cluster. From the same SHOT nearest-neighbor start it lowers mean geodesic error on all six splits, from $0.316/0.417/0.403/0.401/0.231/0.352$ to $0.010/0.035/0.025/0.020/0.020/0.056$ on SCAPE, SCAPE_r, FAUST_r, TOSCA_r, TOPKIDS, and SMAL_r, and it is the lowest stored mean on TOSCA_r (0.020). Pushed labels on SCAPE_r are contiguous on 0.872 of clusters. A FAUST-trained network is 0.013 at home and 0.194 on TOSCA_r. On TOPKIDS, 2D-GEM is 0.006 and LOGOS is 0.020.

1 INTRODUCTION

Take two triangle meshes of one object in two poses. A correspondence is a hard list: every target vertex names the source vertex that should receive its color, motion, or identity. Deformation transfer, performance capture, and statistical body models all consume that list Sumner & Popović (2004); Bouaziz et al. (2013); Anguelov et al. (2005b;a); Loper et al. (2023); Zuffi & Black (2015); Zuffi et al. (2017). Three familiar tools build that list, and they fail in different ways. A short Laplace–Beltrami basis blurs fingers and ears, and a left–right swap that is smooth at low frequency is invisible to what remains Levy (2006); Pinkall & Polthier (1993); Ovsjanikov et al. (2012); Melzi et al. (2019); Rustamov et al. (2013); Donati et al. (2022). Updating each vertex from its neighborhood can match local scores and still deliver one body region as several islands Kamhoua et al. (2022); Kamhoua & Qu (2024). A network trained on one family is accurate at home and weak after a change of category Litany et al. (2017); Sharp et al. (2022); Cao et al. (2023); Eisenberger et al. (2020).

LOGOS (**L**ocal **O**verlap **G**raph **O**ptimization for **S**hape correspondence) edits a hard list that already exists. Three stages do three jobs. A residual gate rewrites only vertices whose neighborhood already disagrees, by greedy multi-hop agreement, and that stage changes the list. Geodesic cleanup then freezes the largest piece of each cluster and renames the islands, and it leaves the list unchanged. A short polish moves points again, only inside the cluster they already belong to. Figure 1 is that pipeline on seven datasets. The experiments ask three questions in that order: whether the editor lowers error from one fixed start, where a network trained on the test family still leads, and whether a low error arrives as one piece.

Contributions.

- A training-free editor with three jobs: residual-gated k-hop reassignment changes $\mathbf{p}$; geodesic cleanup changes only labels; averaging polish changes $\mathbf{p}$ inside a cluster (Sec. 3).
- Four local checks of the three stages, recorded in Appendix B.
- One geodesic-error table of every stored map on SCAPE, SCAPE_r, FAUST_r, TOSCA_r, TOPKIDS, and SMAL_r (Sec. 4). Timings, the knob sweep, and the contiguity counts are in the appendix.

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.

2 Related Work

A functional map stores a correspondence as a small matrix in a Laplace–Beltrami basis, and ZoomOut grows that basis from a coarse map Ovsjanikov et al. (2012); Levy (2006); Pinkall & Polthier (1993); Ren et al. (2018); Rodolí et al. (2017); Rustamov et al. (2013); Donati et al. (2022); Melzi et al. (2019). Fingers, ears, and a smooth left–right swap live in the coefficients that basis drops. NAM refines a hard pointwise map by a neural zoom of the same kind, a linear branch plus a multilayer perceptron, grown across a longer basis Viganò et al. (2025). LOGOS never builds that basis. It edits the vertex list, which is why a spectral zoom and this editor can be run on the same SHOT start and compared as two different repairs. Plain ZoomOut is a stored spectral run whose initializer was not recorded with the row. DIR is a stored orientation method whose geodesic residual assumes a shared vertex order; the class table in Appendix J is the diagnostic of that assumption Xiang et al. (2021).

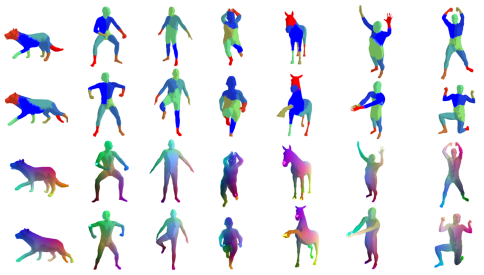

Figure 1: LOGOS on seven datasets, left to right SMAL_r, SCAPE_r, FAUST_r, TOPKIDS, TOSCA_r, FAUST_a, and SCAPE. Top: source clusters. First middle: cleaned clusters warped to the target. Second middle: source map. Bottom: recovered map. FAUST_a is shown here and is not a column of Table 1.

Matching a vertex to a neighborhood that agrees with the source is a quadratic assignment problem Finke et al. (1987); Loiola et al. (2007). 2D-GEM and HOPE update vertices from that local agreement and do not pass the updated map through a cluster step, so one body region can arrive as several islands while the mean geodesic error still looks moderate Kamhoua et al. (2022); Kamhoua & Qu (2024). Appendix I records the schedules those papers publish and sets them next to Table 1: 2D-GEM matches on the 2-hop, and HOPE matches on hops 1 through 8. LOGOS uses a similar k-hop overlap, only on vertices whose residual is already large, and it matches those vertices by a greedy half-approximation of the maximum-weight matching on those scores rather than by the Hungarian algorithm Kuhn (2012). A separate stage then freezes the largest witness piece of each pushed label. SHOT descriptors, and heat and wave kernels, describe a point by its neighborhood and still cannot tell left from right when the neighborhood is mirrored Tombari et al. (2010); Rusu et al. (2009); Sun et al. (2009); Aubry et al. (2011); Kim et al. (2011); Ovsjanikov et al. (2008). SHOT nearest neighbors are the initializer of training-free LOGOS.

Deep functional maps learn basis coefficients from examples Litany et al. (2017); Donati et al. (2020; 2022); Sharp et al. (2022); Cao et al. (2023); Sun et al. (2023). ULRSSM is the unsupervised network in this study, and the suffix on its name is the training set: four networks, not one network scored four ways Cao et al. (2023). On FAUST_r and SCAPE_r, the network trained on that family is the map the editor is measured against. Running LOGOS or NAM on an already strong in-distribution list is a different row of the same table. Surveys separate near-isometry, partial overlap, and topological noise Van Kaick et al. (2011); Tam et al. (2013); Biasotti et al. (2016); Sahillioğlu (2020); Bronstein et al. (2008); Bogo et al. (2014); Anguelov et al. (2005b); Lähner et al. (2016); Cosmo et al. (2016).

3 Method

3.1 Notation

The source mesh is $\mathcal{M}_s = (\mathcal{V}_s, \mathbf{F}_s)$ with coordinates $\mathbf{V}_s \in \mathbb{R}^{n_s \times 3}$; the target is $\mathcal{M}_t$ with $\mathbf{V}_t \in \mathbb{R}^{n_t \times 3}$. The hard map $\mathbf{p} : \mathcal{V}_t \to \mathcal{V}_s$ stores, at position j, the source index $\mathbf{p}(j)$. The transport $\Pi \in \{0,1\}^{n_t \times n_s}$ has $\Pi_{j,\mathbf{p}(j)} = 1$, so $\Pi \widetilde{\mathbf{A}}_s \Pi^\top$ pulls the source adjacency onto target indices. $\mathbf{A}$ is the mesh 1-skeleton and $\widetilde{\mathbf{A}}$ is the witness copy below. $\mathbf{D}$ is the diagonal degree matrix of the matrix it multiplies. $\mathbf{H}^{(k)}$ marks pairs with a walk of length k and none of length $k-1$; on a clean triangle mesh that set equals graph distance k, and the runs build $\mathbf{H}^{(k)}$ on the witness graph. Cluster labels

$\mathbf{c}_s$ and $\mathbf{c}_t$ are geodesic Voronoi indices, not semantic parts. $\mathcal{N}_t(j)$ is the set of target vertices joined to j by a positive witness weight. d_g is the exact geodesic length in the appendix propositions, and d_{heat} is the heat-method distance in the reported error Crane et al. (2013). Appendix A lists the same symbols.

3.2 Where the hard map comes from

The refiner consumes a target-to-source index array from one of three places: SHOT nearest neighbors Tombari et al. (2010), which is training-free LOGOS; landmarks warped by the ground-truth map, which is LOGOS-lmks; or a ULRSSM network trained on one dataset Cao et al. (2023), which is LOGOS-ULRSSM-*. NAM-SHOT and the NAM-ULRSSM rows edit those named starts. No NAM landmark row is stored. Figure 2 is that SHOT-initialized comparison on SCAPE_r, TOPKIDS, SMAL_r, and TOSCA_r. Training-free, in this paper, means no dataset training, and it names LOGOS. NAM fits a network at test time.



Figure 2: SHOT initialization. Colors are source coordinates transferred to the target. Within each dataset, left to right: source, LOGOS, ZoomOut, 2D-GEM, HOPE, and DIR. Top-left: SCAPE_r. Top-right: TOPKIDS. Middle: SMAL_r. Bottom: TOSCA_r.

3.3 Witness graph

Both meshes draw from one generator, seed 42. The draw flips 1% of the edges that bound two faces, deletes three walks of 5 to 11 edges, adds $\lfloor 0.005\, n \rfloor$ chords between vertices at walk distance 2 (n is the vertex count of that mesh), and multiplies each directed edge by an independent weight in $[0.85, 1.15]$. Hop shells use the binarized support, so the weights do not change who is in $\mathbf{H}^{(k)}$. The residual uses the weighted copy. Geodesics stay on the original mesh. Every stored LOGOS map uses this perturbation, so the table is one protocol. On SCAPE the two meshes share connectivity and the ground truth is the identity, so the shared seed can make the witness graphs agree in index space. Remeshed pairs are the informative test of the switch.

3.4 Who is allowed to move

Farthest-point sampling places $K{=}10$ seeds on the source and paints $\mathbf{c}_s$ by the nearest seed. At iteration $t = 1, \ldots, T$ with $T{=}32$, the hop is $k = 1 + ((t-1) \bmod 8)$. The residual of a target vertex is the relative ℓ_1 gap between its witness row and the row the current map transports from the source:

$$\varepsilon(j) = \frac{\left\| (\Pi \widetilde{\mathbf{A}}_s \Pi^\top - \widetilde{\mathbf{A}}_t)[j,:] \right\|_1}{\|\widetilde{\mathbf{A}}_t[j,:]\|_1 + \delta}, \tag{1}$$

with $\delta = 10^{-8}$ only so that a zero target row does not divide by zero. The set $\mathcal{B} = \{j : \varepsilon(j) \geq \tau_{\text{fix}}\}$ uses $\tau_{\text{fix}}{=}0.95$, so only the tail of the residual is rewritten and a calm neighborhood stays put. If the transported witness row equals the target row, then $\varepsilon(j) = 0$. The score of sending $j \in \mathcal{B}$ to source vertex i reindexes the source shell by the current map:

$$s_k(i,j) = \mathbf{H}_s^{(k)} \left[i, \mathbf{p}(\cdot) \right]^\top \mathbf{H}_t^{(k)}[:, j]. \tag{2}$$

The score s_k is not the residual. The residual asks whether the current list already reproduces the target's witness row. The score asks, among vertices that failed that test, which source index has a k-hop neighborhood overlapping the target vertex once both neighborhoods are written in the current map. A high residual grants permission to move. A high score chooses the destination. Edges are then taken in decreasing score, and an edge that touches a used vertex is rejected. That scan is a $\frac{1}{2}$-approximation of the maximum-weight matching on those scores (Appendix B). The ratio is about score weight on that iteration. An unmatched all-zero column is set to source index 0 and can collide, so it is outside the $\frac{1}{2}$ score-weight guarantee.

3.5 LABELS, THEN A POLISH INSIDE THE CLUSTER

Push the source labels across the current map, $\mathbf{c}_t = \mathbf{c}_s \circ \mathbf{p}$. Cleanup freezes the largest witness component of each label, relabels the other vertices of that label by the nearest frozen anchor, and does not write $\mathbf{p}$; contiguity of $\mathbf{c}_s \circ \mathbf{p}$ is a property of the finished list. The argument is Proposition 2, and it uses exact length.

For $\lceil T/2 \rceil$ iterations, each iteration sets $\mathbf{x} = \mathbf{V}_s \circ \mathbf{p}$ and $\mathbf{y} = \mathbf{M}_t\mathbf{x}$ with $\mathbf{M}_t = \mathbf{D}^{-1}\widetilde{\mathbf{A}}_t$, one multiplication, and does not overwrite $\mathbf{x}$. Let $\rho(j) = \|\mathbf{x}(j) - \mathbf{y}(j)\|_2$. The value q_{lap}=0.95 is a quantile of ρ, not of ε. The snap sends each vertex at or above that quantile to

$$i^\star = \arg \min_{i:\, \mathbf{c}_s(i) = \mathbf{c}_t(j)} \big\|\mathbf{y}(j) - \mathbf{V}_s(i)\big\|_2 \tag{3}$$

inside its cluster. The distance is ambient Euclidean distance in the source coordinate frame. Each snap is independent across j, so the finished map need not be a bijection. One symmetrized step with positive degrees does not raise Dirichlet energy; the reported walk is the directed operator $\mathbf{M}_t$, so that inequality is Proposition 3 and does not cover this step.

Algorithm 1: LOGOS. $\mathcal{T}$ perturbs witness adjacency. Reported runs use $k_{\max}$=8.

Input: Meshes $\mathcal{M}_s, \mathcal{M}_t$; hard map $\mathbf{p}^{(0)}$; $K, k_{\max}, T, \tau_{\text{fix}}, q_{\text{lap}}$
Output: Hard map $\mathbf{p}$

$\widetilde{\mathbf{A}}_s, \widetilde{\mathbf{A}}_t \leftarrow \mathcal{T}(\mathbf{A}_s), \mathcal{T}(\mathbf{A}_t)$;
$\mathbf{c}_s \leftarrow$ geodesic Voronoi labels on $\mathcal{M}_s$;
$\mathbf{p} \leftarrow \mathbf{p}^{(0)}$;
for $t \leftarrow 1$ **to** T **do**
  $k \leftarrow ((t-1) \bmod k_{\max}) + 1$;
  $\mathcal{B} \leftarrow \{j : \varepsilon(j) \geq \tau_{\text{fix}}\}$ from (1);
  if $\mathcal{B} = \emptyset$ **then**
    break
  Reassign $\mathbf{p}|_{\mathcal{B}}$ by greedy matching on s_k (2);
  An unmatched all-zero column is set to source index 0;
Update labels $\mathbf{c}_t$ by largest-component cleanup ($\mathbf{p}$ unchanged);
for $t \leftarrow 1$ **to** $\lceil T/2 \rceil$ **do**
  $\mathbf{x} \leftarrow \mathbf{V}_s \circ \mathbf{p}$; $\mathbf{y} \leftarrow \mathbf{M}_t\mathbf{x}$ with $\mathbf{M}_t = \mathbf{D}^{-1}\widetilde{\mathbf{A}}_t$;
  $\rho(j) \leftarrow \|\mathbf{x}(j) - \mathbf{y}(j)\|_2$;
  Reassign only j with $\rho(j)$ at or above the q_{lap} quantile of ρ, by (3);
return $\mathbf{p}$

3.6 WHY THE STAGES ARE SEPARATE

The three failures in the introduction are three different objects, and each stage returns a different one.

A short Laplace–Beltrami basis cannot see a swap that is already smooth at low frequency. The editor does not build that basis. It rewrites a vertex list, and the residual in (1) decides where. If the transported witness row equals the target row, the residual is zero and the greedy stage leaves that source index in place. A calm neighborhood is evidence the current list already agrees locally, which is why a strong initializer is worth keeping and why the gate spends its edits on the tail.

Neighborhood scores can still split one body region into islands. The greedy stage does not try to repair that. Cleanup relabels those islands and leaves $\mathbf{p}$ unchanged. It freezes the largest witness component of each pushed label and renames every other vertex of that label to the nearest frozen anchor. Algorithm 1 writes that update into $\mathbf{c}_t$. The index array $\mathbf{p}$ waits for the polish. Contiguity, measured later, is a property of the list the polish returns.

The polish is the stage that moves indices again, and it moves them inside a cluster. Source coordinates are pulled onto the target by the current list, $\mathbf{x} = \mathbf{V}_s \circ \mathbf{p}$, the directed walk $\mathbf{y} = \mathbf{M}_t\mathbf{x}$ is one multiplication, and the snap keeps vertices at or above the q_{lap} quantile of $\rho = \|\mathbf{x} - \mathbf{y}\|_2$. The snap

in (3) compares the smoothed point with source vertices that already carry the same cluster name, in the source coordinate frame. A finger is not free to jump to a toe. Each snap is independent, so the finished list need not be bijective.

Proposition 1 checks row equality, Proposition 4 checks score weight, Proposition 2 checks exact-length anchors, and Proposition 3 checks one symmetrized step. The runs use the heat method, witness-graph hops, and the directed walk $\mathbf{M}_t$.

4 Experiments

4.1 What is measured

For every target vertex that has a ground-truth source, the error is the heat-method distance on the target between that vertex and the target vertex the ground truth assigns to the predicted source, divided by $\sqrt{\mathrm{Area}(\mathcal{M}_t)}$:

$$E(\mathbf{p}) = \mathrm{mean}_j \, \frac{d_{\mathrm{heat}}\big(j,\, t^\star(\mathbf{p}(j))\big)}{\sqrt{\mathrm{Area}(\mathcal{M}_t)}}. \tag{4}$$

Here $t^\star$ sends a source vertex to its ground-truth target vertex. A value near 0.05 is about five percent of a typical length of the shape. The splits are the pair lists of this study: SCAPE (71), SCAPE_r (71), FAUST_r (300), TOSCA_r (284), TOPKIDS (90), and SMAL_r (298). Ground truth is the identity on SCAPE, a shared reference index on SCAPE_r, FAUST_r, and TOSCA_r, the intersection of the reference maps on TOPKIDS, and a column-by-column pairing on SMAL_r. The same rule is used for every method (Appendix E). Shared settings are $K{=}10$, hops 1–8, $T{=}32$, witness seed 42, $\tau_{\mathrm{fix}}{=}0.95$, and $q_{\mathrm{lap}}{=}0.95$.

4.2 One table

Table 1 is every stored geodesic error. Teal is the lowest entry in a column and orange is the second lowest. The three questions in the introduction are three readings of this one matrix: the SHOT, NAM-SHOT, and LOGOS rows; the network rows and their edited copies; and, in the appendix, the contiguity of the list those errors come from.

Table 1: Mean geodesic error (↓) on the pair lists of this study. The lists are longer than several published protocols, and every column uses that same list. Teal is lowest and orange is second. A dash was not evaluated. Plain ZoomOut's initializer was not recorded. SCAPE shares connectivity and identity ground truth.

Method	SCAPE	SCAPE_r	FAUST_r	TOSCA_r	TOPKIDS	SMAL_r
SHOT k-NN	0.316	0.417	0.403	0.401	0.231	0.352
NAM-SHOT	0.061	0.167	0.136	0.130	0.048	0.140
NAM-ULRSSM-faust	0.058	0.077	**0.015**	0.090	0.337	–
NAM-ULRSSM-scape	0.016	**0.027**	0.030	0.082	0.310	–
NAM-ULRSSM-topkids	–	0.234	0.176	–	0.134	–
NAM-ULRSSM-smal	–	–	–	–	–	**0.025**
ULRSSM-faust	0.154	0.093	**0.013**	0.194	0.354	–
ULRSSM-scape	0.041	**0.024**	0.064	0.207	0.332	0.184
ULRSSM-topkids	0.228	0.230	0.166	0.222	0.137	0.242
ULRSSM-smal	–	–	–	–	–	0.051
ZoomOut	0.026	0.074	0.048	0.059	0.104	0.133
2D-GEM	0.014	0.097	0.087	0.085	**0.006**	0.187
HOPE	**0.011**	0.049	0.039	**0.030**	0.029	0.077
ZoomOut-lmks	0.059	0.066	0.043	0.111	–	0.312
ZoomOut-ULRSSM-faust	0.111	0.094	0.034	0.180	–	–
ZoomOut-ULRSSM-scape	0.054	0.041	0.057	0.185	–	0.133
ZoomOut-ULRSSM-topkids	–	–	–	–	0.159	–
ZoomOut-ULRSSM-smal	–	–	–	–	–	0.032
LOGOS-lmks	0.044	0.048	0.028	0.061	0.072	0.033
LOGOS-ULRSSM-faust	0.074	0.069	0.016	0.066	–	–
LOGOS-ULRSSM-scape	0.018	**0.027**	0.022	0.063	0.154	0.108
LOGOS-ULRSSM-topkids	–	–	–	–	0.111	–
LOGOS-ULRSSM-smal	–	–	–	–	–	**0.028**
LOGOS	**0.010**	0.035	0.025	**0.020**	**0.020**	0.056

4.3 From one hard start

SHOT nearest neighbors, NAM-SHOT, and training-free LOGOS share an initializer. NAM-SHOT fits a test-time zoom to that list and lowers the mean on every split. LOGOS lowers it further on every split, from 0.316/0.417/0.403/0.401/0.231/0.352 to 0.010/0.035/0.025/0.020/0.020/0.056. The middle row is 0.061/0.167/0.136/0.130/0.048/0.140, so the drop is visible against a spectral editor of the same start, not only against the descriptor match.

On TOSCA_r the stored matrix has no TOSCA-trained network, and LOGOS at 0.020 is the lowest mean, with HOPE next at 0.030. The sample standard deviations of the six training-free pair means are 0.005/0.044/0.023/0.013/0.013/0.059. The TOSCA_r gap of 0.010 and the TOSCA_r spread of 0.013 are the same order. On original SCAPE, which shares connectivity and uses the identity as ground truth, LOGOS at 0.010 is the lowest entry and sits next to HOPE at 0.011, inside a pair spread of 0.005. The landmark start of the same editor is 0.044 on that check. Landmarks are a different initializer, and on a shared-connectivity pair they are not the lower one.

The remeshed human column is the place to read the whole matrix at once. On SCAPE_r the SCAPE-trained network is 0.024, both stored refinements of that list are 0.027, and training-free LOGOS is 0.035. HOPE is 0.049, ZoomOut is 0.074, 2D-GEM is 0.097, and SHOT is 0.417. A FAUST-trained network on this column is 0.093, and LOGOS on that list reaches 0.069, still above the training-free editor. A TOPKIDS-trained network is 0.230. Matched training wins this column. A foreign training set does not. The training-free editor is the lowest map in the column that did not train on SCAPE.

4.4 Trained maps, at home and after a change of category

On FAUST_r and SCAPE_r the network trained on that family has the lowest stored mean. The FAUST-trained network is 0.013; NAM-ULRSSM-faust is 0.015 and LOGOS-ULRSSM-faust is 0.016. The network's own list is the one to keep there. Training-free LOGOS on FAUST_r is 0.025. The SCAPE-trained network is 0.024 on SCAPE_r, and both stored refinements of that list sit at 0.027.

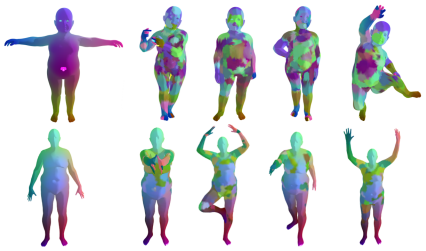

Figure 3: Train/test shift, the second reading of Table 1. Top: a FAUST-trained network on TOPKIDS. Bottom: a TOPKIDS-trained network on FAUST_r.

The same FAUST-trained network is 0.194 on TOSCA_r. Editing that list with LOGOS brings it to 0.066, and the SHOT-initialized editor is lower still at 0.020. The SCAPE-trained network is 0.207 on TOSCA_r; LOGOS-ULRSSM-scape is 0.063 there. Off the training family, the editor is a repair of the network's list, and the training-free start is the lowest of those TOSCA_r entries. Figure 3 is the same shift as color. Table 2 records the four-column version: a FAUST-trained network at 0.354 on TOPKIDS against training-free LOGOS at 0.020, and a TOPKIDS-trained network at 0.166 on FAUST_r against 0.025.

Table 2: Geodesic error (↓) on four test sets. The ULRSSM header is that network's training set. The LOGOS row is the training-free SHOT-initialized refiner, so the two FAUST_r entries match.

Test set ULRSSM train set	FAUST_r FAUST	SCAPE_r SCAPE	TOPKIDS FAUST	FAUST_r TOPKIDS
ULRSSM	0.013	0.024	0.354	0.166
LOGOS	0.025	0.035	0.020	0.025

TOPKIDS is scored on the intersection of the reference maps. 2D-GEM is 0.006 and LOGOS is 0.020. The network trained on TOPKIDS is 0.137, and LOGOS on that list is 0.111. On this split the local structural method leads, the training-free editor is next, and the network trained on the test set does not lead.

On SMAL_r the lowest stored map is NAM-ULRSSM-smal at 0.025, then LOGOS-ULRSSM-smal at 0.028 and ZoomOut-ULRSSM-smal at 0.032, against the unedited SMAL network at 0.051. The trained list is the start the refiners improve. Among maps that were run on every split, LOGOS-lmks is 0.033 and training-free LOGOS is 0.056, so landmarks are the lower of those two on the animal set and the higher of the two on the other five.

Remeshing reorders the same names. ULRSSM-scape is 0.041 on original SCAPE and 0.024 on SCAPE_r. Training-free LOGOS is 0.010 and 0.035 on those two columns. A change of tessellation is its own column, and the witness switch is informative on the remeshed one, where the source and target no longer share an index.

4.5 Error in one piece

Deformation transfer consumes a region. A small mean can still paint one limb as many islands, which is the second failure in the introduction. Table 3 measures the pushed labels $\mathbf{c}_s \circ \mathbf{p}$ on the clean target edges, after the full pipeline. On SCAPE_r, LOGOS averages 1.25 ± 0.30 components per cluster and is contiguous on 0.872 of clusters, against 0.324 for HOPE, 0.172 for ZoomOut, and 0.044 for 2D-GEM. On that split 0.201 of mapped vertices share a source index, so the contiguous list still repeats some source vertices. HOPE, ZoomOut, and 2D-GEM have collision 0 there and many more pieces. On TOPKIDS the contiguous fractions are 0.602 for LOGOS, 0.189 for 2D-GEM, 0.106 for HOPE, and 0.007 for ZoomOut. 2D-GEM has the lower TOPKIDS geodesic error and the more fragmented labels. The two measurements answer different questions. On ten SCAPE_r pairs the same labels fall from 145.4 components per cluster under SHOT, with no contiguous cluster, to 1.17 components and a contiguous fraction of 0.91 under the full editor (Appendix H).

Table 3: Pushed-label contiguity of $\mathbf{c}_s \circ \mathbf{p}$ on the clean target edges, after the full pipeline. Cleanup does not write $\mathbf{p}$. Collision is the fraction of mapped vertices that share a source index.

Dataset	Map	Pairs	CC / cluster	Contiguous	Collision
SCAPE_r	LOGOS	71	1.25 ± 0.30	0.872	0.201
SCAPE_r	HOPE	71	5.29 ± 2.67	0.324	0.000
SCAPE_r	ZoomOut	71	8.10 ± 2.47	0.172	0.000
SCAPE_r	2D-GEM	71	17.36 ± 2.94	0.044	0.000
TOPKIDS	LOGOS	90	2.31 ± 0.96	0.602	0.185
TOPKIDS	2D-GEM	90	7.07 ± 2.70	0.189	0.000
TOPKIDS	HOPE	90	41.5 ± 35.4	0.106	0.061
TOPKIDS	ZoomOut	90	90.7 ± 74.3	0.007	0.231

4.6 Which knob moves the error

Appendix Table 7 changes one setting on ten pairs per dataset. Opening the residual gate raises that ten-pair mean on every set except original SCAPE, where it stays 0.013. On the full SCAPE_r file the same removal moves 0.035 ± 0.045 to 0.040 ± 0.041 and the refiner-only time from 1.35 s to 2.00 s (Table 4). The 0.045 is that file's standard deviation, not the 0.044 pair standard deviation of the main SCAPE_r row. The gate is the part that keeps a calm vertex put.

Table 4: Leave-one-stage ablations on the full SCAPE_r split (n=71). Mean±std of pair-wise mean geodesic error. Seconds are refiner-only.

Setting	Mean±std	Seconds / pair
Default	0.035 ± 0.045	1.35
No perturbation	0.033 ± 0.039	1.19
No residual gate	0.040 ± 0.041	2.00
No component cleanup	0.037 ± 0.046	1.19
No averaging	0.041 ± 0.045	1.18

The polish quantile is the sensitive knob on the identity check. Setting $q_{\text{lap}}{=}0.50$ raises the ten-pair SCAPE mean from 0.013 to 0.060; $q_{\text{lap}}{=}0.99$ is 0.009 there. Snapping half the vertices by displacement is too wide when the ground truth is the identity. The reported runs use 0.95. Removing the polish on the full SCAPE_r file raises the mean from 0.035 ± 0.045 to 0.041 ± 0.045, the largest leave-one-stage move in that file. On the ten-pair identity check, removing it lowers the mean from 0.013 to 0.009. The operating point matters more than the presence of the loop.

Hops 1–4 raise the ten-pair mean on five of the six columns; TOPKIDS stays 0.019. Hops 1–12 are the lowest SMAL_r entry in the grid, 0.058 against the ten-pair default 0.078. The reported protocol is hops 1–8 on every split, including SMAL_r. The cluster count is a flat direction of the same grid. $K{=}5$ and $K{=}20$ stay within a few thousandths of $K{=}10$ on the human columns. $K{=}20$ is slightly lower on TOPKIDS (0.017 against 0.019) and on SMAL_r (0.074 against 0.078). The reported choice $K{=}10$ sits in the middle of that range, and the polish cannot cross whatever clusters that choice paints. The residual threshold is more directional. $\tau_{\text{fix}}{=}0.50$ raises the ten-pair means on SCAPE_r, TOSCA_r, and SMAL_r (0.025, 0.026, 0.089 against 0.022, 0.023, 0.078). $\tau_{\text{fix}}{=}0.99$ is level with the default on the human columns and slightly lower on SMAL_r (0.073). The reported runs use 0.95, between those two settings. Turning the witness perturbation off leaves the ten-pair human means almost unchanged and raises SMAL_r from 0.078 to 0.097. On the full SCAPE_r file the same removal is 0.033 ± 0.039 against 0.035 ± 0.045. The perturbation changes which ties break. On the human ten-pair means the removal is almost flat, and on the animal cut the perturbed protocol is the lower one. Cleanup's geodesic movement is small on both summaries: 0.037 ± 0.046 against 0.035 ± 0.045 on the full SCAPE_r file, and mixed thousandths on the ten-pair grid. That fits a stage whose output is a label field. Changing only the witness seed on the ten SCAPE_r pairs gives 0.024, 0.022, 0.022, 0.022, and 0.022 for seeds 0, 1, 2, 42, and 123.

4.7 Which row to take

The matrix is easier to use as a choice than as a ranking. Four situations cover the columns.

When the two meshes share connectivity and no training set is available, the training-free editor is the lowest stored mean. That is original SCAPE, at 0.010, next to HOPE at 0.011. The landmark start of the same editor is 0.044 on that check, so landmarks are not a substitute for the SHOT start when the index is already shared.

When the test family is one the network was trained on, and the meshes have been remeshed, keep the network's list. That is FAUST_r at 0.013 and SCAPE_r at 0.024. Both stored edits of those lists sit a few thousandths higher. The editor is not a replacement for a matched training set.

When the test family is not the training family, the editor is the repair. On TOSCA_r, where no TOSCA-trained network is stored, training-free LOGOS at 0.020 is the lowest mean, and it also brings a FAUST-trained list from 0.194 to 0.066. On TOPKIDS the goal decides the row. 2D-GEM has the lower geodesic error, 0.006 against 0.020, and LOGOS has the more contiguous labels, 0.602 against 0.189. The network trained on TOPKIDS, 0.137, is neither.

When the shapes are animals, start from the SMAL-trained list and refine it. NAM-ULRSSM-smal at 0.025 is the lowest stored mean, and LOGOS on that same list is 0.028. If the start has to be training-free, landmarks at 0.033 are lower than SHOT-initialized LOGOS at 0.056, which is the reverse of the shared-connectivity check.

4.8 Scope

The hop score counts overlap, so an intrinsic symmetry is another high-overlap list. The witness perturbation changes which ties break. Partial overlap and a change of genus are outside the edit Rodolí et al. (2017); Lähner et al. (2016). TOPKIDS does not test vertices that the reference maps leave unmatched. The fractions of training-free pairs above 0.1 are $0/0.070/0.030/0.011/0/0.225$, and above 0.15 are $0/0.056/0.003/0/0/0.077$. The tail is a spread of geodesic error. It is empty on SCAPE and TOPKIDS and largest on SMAL_r.

4.9 FIGURES AND TIMINGS

The three figures are the same three questions, drawn rather than tabulated. Figure 1 is the same pipeline on seven datasets, left to right SMAL_r, SCAPE_r, FAUST_r, TOPKIDS, TOSCA_r, FAUST_a, and SCAPE: source clusters, cleaned clusters on the target, the source map, and the recovered map. Figure 2 is the SHOT-initialized comparison with ZoomOut, 2D-GEM, HOPE, and DIR on SCAPE_r, TOPKIDS, SMAL_r, and TOSCA_r. Figure 3 is that category change on a pair: a FAUST-trained network on TOPKIDS, and a TOPKIDS-trained network on FAUST_r. Appendix Figures 4 and 5 continue that color test on SCAPE and on a network start. They are the same editor.

Wall-clock times are Appendix Table 6. Two machines produced them, so the column is a record of the run that wrote each row, not a ranking. On the machine that timed the editor, training-free LOGOS takes $6.9\,\mathrm{s}$ per FAUST_r pair and $6.7\,\mathrm{s}$ on TOSCA_r, and NAM-SHOT takes $24.1\,\mathrm{s}$ and $24.2\,\mathrm{s}$ on those splits. Those four cells are a record of that machine. The geodesic table is the comparison. Teal and orange in Table 1 are decided inside each column. A dash is a map that was not run on that split. On SCAPE_r the two refinements of the SCAPE-trained list tie for second at 0.027, and both cells are marked. The appendix records one argument for each displayed proposition, the wall-clock times, and the contiguity counts cited above. The same three stages produce every LOGOS row in the table. What changes from row to row is the initializer. The per-class orientation diagnostic, the ten-pair lists, and the witness-seed table follow the proofs, so the main text can stay with this one matrix.

5 CONCLUSION

LOGOS edits a hard map in three stages, and the stages answer three different questions. From SHOT nearest neighbors the edit lowers geodesic error on all six splits, and on TOSCA_r the mean 0.020 is the lowest stored entry. On FAUST_r and SCAPE_r the network trained on that family remains the list to keep, at 0.013 and 0.024; off that family the same editor repairs the network's list, and the training-free start is lower still on TOSCA_r. On SCAPE_r the finished labels are contiguous on 0.872 of clusters. The editor is the map to run from a SHOT start when the test family has no matched network, and it is the repair to run when a network has left its training family. Where a matched network already leads, the table says to keep that list. The contiguity counts are what make those training-free maps usable as regions, rather than as a mean alone.

AI USE STATEMENT

Generative tools were used to edit the manuscript and to check that tabulated numbers match stored measurements. They were not used to invent measurements, proofs, or citations. The authors reviewed the text, proofs, and experimental claims and take responsibility for the paper.

ETHICS STATEMENT

Shape correspondence methods can be applied in animation, AR/VR, and medical geometry. Training-free refiners reduce dependence on proprietary scans; they can also be misused in biometric matching. This work targets public geometry-processing benchmarks.

REPRODUCIBILITY STATEMENT

Datasets, pair counts, and correspondence conventions are stated in Sec. 4.1 and Appendix E. Shared refine settings are $K{=}10$, hops 1–8, $T{=}32$, witness seed 42, absolute residual threshold $\tau_{\mathrm{fix}}{=}0.95$, and polish displacement quantile $q_{\mathrm{lap}}{=}0.95$. NAM uses Neural ZoomOut ($k_{\mathrm{ini}}{=}20$, $k_{\mathrm{end}}{=}200$, step 20; MLP width 128; Adam with early stopping). NAM-SHOT starts from SHOT k-NN; NAM-ULRSSM-* starts from the ULRSSM network trained on that set, under that method's native mesh normalization. Geodesic error is the area-normalized heat method on the target mesh. Ten-pair ablation subsets are listed in Table 9.

REFERENCES

Dragomir Anguelov, Praveen Srinivasan, Hoi-Cheung Pang, Daphne Koller, Sebastian Thrun, and James Davis. The correlated correspondence algorithm for unsupervised registration of nonrigid surfaces. In *Advances in neural information processing systems*, pp. 33–40, 2005a.

Dragomir Anguelov et al. Scape: shape completion and animation of people. *ACM Transactions on Graphics (TOG)*, 2005b.

Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pp. 1626–1633, 2011. doi: 10.1109/ICCVW.2011.6130444.

Silvia Biasotti, Andrea Cerri, Alex Bronstein, and Michael Bronstein. Recent trends, applications, and perspectives in 3d shape similarity assessment. *Computer Graphics Forum*, 35(6):87–119, 2016.

Federica Bogo, Javier Romero, Matthew Loper, and Michael J. Black. Faust: Dataset and evaluation for 3d mesh registration. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3794–3801, 2014. doi: 10.1109/CVPR.2014.491.

Sofien Bouaziz, Yangang Wang, and Mark Pauly. Online modeling for realtime facial animation. *ACM Transactions on Graphics (TOG)*, 2013.

Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. Numerical geometry of non-rigid shapes. *Springer Science & Business Media*, 2008.

Dongliang Cao, Paul Roetzer, and Florian Bernard. Unsupervised learning of robust spectral shape matching. *ACM Transactions on Graphics (TOG)*, 42(4):1–15, 2023.

Luca Cosmo, Emanuele Rodolà, Michael M Bronstein, Andrea Torsello, Daniel Cremers, and Yusuf Sahillioglu. Shrec'16: Partial shape matching of non-rigid objects. In *Eurographics Workshop on 3D Object Retrieval (3DOR)*, 2016.

Keenan Crane et al. Geodesics in heat. *ACM Transactions on Graphics (TOG)*, 2013.

Nicolas Donati, Etienne Corman, and Maks Ovsjanikov. Deep Orientation-Aware Functional Maps: Tackling Symmetry Issues in Shape Matching. *CVPR "12 : IEEE Conference on Computer Vision and Pattern Recognition*, 2022. doi: 10.48550/arXiv.2204.13453. URL `https://hal.archives-ouvertes.fr/hal-03654308`.

Nicolas Donati et al. Deep geometric functional maps. *CVPR*, 2020.

Marvin Eisenberger, Zorah Lähner, and Daniel Cremers. Smooth shells: Multi-scale shape registration with functional maps. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 12262–12271, 2020. doi: 10.1109/CVPR42600.2020.01228.

Aoxiang Fan, Jiayi Ma, Xin Tian, Xiaoguang Mei, and Wei Liu. Coherent point drift revisited for non-rigid shape matching and registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1424–1434, June 2022.

Gerd Finke, Rainer E. Burkard, and Franz Rendl. Quadratic assignment problems. In Silvano Martello, Gilbert Laporte, Michel Minoux, and Celso Ribeiro (eds.), *Surveys in Combinatorial Optimization*, volume 132 of *North-Holland Mathematics Studies*, pp. 61–82. North-Holland, 1987. doi: https://doi.org/10.1016/S0304-0208(08)73232-8.

Barakeel Fanseu Kamhoua and Huamin Qu. HOPE: Shape matching via aligning different k-hop neighbourhoods. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

Barakeel Fanseu Kamhoua, Lin Zhang, Yongqiang Chen, Han Yang, Kaili Ma, Bo Han, Bo Li, and James Cheng. Exact shape correspondence via 2d graph convolution. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

Vladimir G Kim et al. Blended intrinsic maps. *ACM Transactions on Graphics (TOG)*, 2011.

H. Kuhn. The hungarian method for the assignment problem. *Naval Research Logistic Quarterly*, 2, 05 2012.

Z. Lähner, E. Rodolà, M. M. Bronstein, D. Cremers, O. Burghard, L. Cosmo, A. Dieckmann, R. Klein, and Y. Sahillioğlu. Matching of deformable shapes with topological noise. In *Proceedings of the Eurographics 2016 Workshop on 3D Object Retrieval*, 3DOR '16, pp. 55–60, Goslar, DEU, 2016. Eurographics Association. ISBN 9783038680048.

Zorah Lähner, Emanuele Rodolà, Michael M Bronstein, Denis Stoll, and Daniel Cremers. Shrec'16: Matching of deformable shapes with topological variations. In *Eurographics Workshop on 3D Object Retrieval (3DOR)*, 2016.

Bruno Levy. Laplace-beltrami eigenfunctions for deformation and editing. *SGP*, 2006.

Or Litany, Tal Remez, Emanuele Rodola, Alex Bronstein, and Michael Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Oct 2017.

Eliane Loiola, Nair Abreu, Paulo Boaventura-Netto, Peter Hahn, and Tania Querido. A survey of the quadratic assignment problem. *European Journal of Operational Research*, 176:657–690, 01 2007. doi: 10.1016/j.ejor.2005.09.032.

Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. *SMPL: A Skinned Multi-Person Linear Model*. Association for Computing Machinery, New York, NY, USA, 1 edition, 2023. ISBN 9798400708978. URL https://doi.org/10.1145/3596711.3596800.

Cheng Mao, Mark Rudelson, and Konstantin Tikhomirov. Exact matching of random graphs with constant correlation. *Probability Theory and Related Fields*, 186(1-2):327–389, 2023.

Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. Zoomout: spectral upsampling for efficient shape correspondence. *ACM Trans. Graph.*, 38(6), November 2019. ISSN 0730-0301. doi: 10.1145/3355089.3356524. URL https://doi.org/10.1145/3355089.3356524.

Maks Ovsjanikov, Jian Sun, and Leo Guibas. Global intrinsic symmetries of shapes. *Comp. Graph. Forum*, 27(5):1341–1348, 2008.

Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.*, 31(4), July 2012. ISSN 0730-0301. doi: 10.1145/2185520.2185526. URL https://doi.org/10.1145/2185520.2185526.

Ulrich Pinkall and Konrad Polthier. Computing Discrete Minimal Surfaces and their Conjugates. *Experimental mathematics*, 2(1):15–36, 1993.

Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. *ACM Trans. Graph.*, 37(6), December 2018. ISSN 0730-0301. doi: 10.1145/3272127.3275040. URL https://doi.org/10.1145/3272127.3275040.

E. Rodolí, L. Cosmo, M. M. Bronstein, A. Torsello, and D. Cremers. Partial functional correspondence. *Comput. Graph. Forum*, 36(1):222–236, January 2017. ISSN 0167-7055. doi: 10.1111/cgf.12797. URL https://doi.org/10.1111/cgf.12797.

Raif M. Rustamov, Maks Ovsjanikov, Omri Azencot, Mirela Ben-Chen, Frédéric Chazal, and Leonidas Guibas. Map-based exploration of intrinsic shape differences and variability. *ACM Trans. Graph.*, 32(4), July 2013. ISSN 0730-0301. doi: 10.1145/2461912.2461959. URL https://doi.org/10.1145/2461912.2461959.

Radu Bogdan Rusu, Nico Blodow, and Michael Beetz. Fast point feature histograms (fpfh) for 3d registration. In *2009 IEEE International Conference on Robotics and Automation*, pp. 3212–3217, 2009. doi: 10.1109/ROBOT.2009.5152473.

Yusuf Sahillioğlu. Recent advances in shape correspondence. *Vis. Comput.*, 36(8):1705–1721, August 2020. ISSN 0178-2789. doi: 10.1007/s00371-019-01760-0. URL https://doi.org/10.1007/s00371-019-01760-0.

Nicholas Sharp, Souhaib Attaiki, Keenan Crane, and Maks Ovsjanikov. Diffusionnet: Discretization agnostic learning on surfaces. *ACM Transactions on Graphics (TOG)*, 41(3):1–16, 2022.

Robert W Sumner and Jovan Popović. Deformation transfer for triangle meshes. *ACM Transactions on Graphics (TOG)*, 23(3):399–405, 2004.

Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Proceedings of the Symposium on Geometry Processing*, SGP '09, pp. 1383–1392, Goslar, DEU, 2009. Eurographics Association.

Mingze Sun, Shiwei Mao, Puhua Jiang, Maks Ovsjanikov, and Ruqi Huang. Spatially and spectrally consistent deep functional maps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 14497–14507, 2023.

Gary KL Tam, Zhi-Quan Cheng, Yu-Kun Lai, Frank C Langbein, Yonghuai Liu, David Marshall, Ralph R Martin, Xian-Fang Sun, and Paul L Rosin. Registration of 3D point clouds and meshes: a survey from rigid to nonrigid. *IEEE TVCG*, 19(7):1199–1217, 2013.

Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique signatures of histograms for local surface description. In *Proceedings of the 11th European Conference on Computer Vision Conference on Computer Vision: Part III*, ECCV'10, pp. 356–369, Berlin, Heidelberg, 2010. Springer-Verlag. ISBN 364215557X.

Oliver Van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 30(6):1681–1707, 2011.

Giulio Viganò, Maks Ovsjanikov, and Simone Melzi. NAM: Neural adjoint maps for refining shape correspondences. *ACM Trans. Graph.*, 44(4), July 2025. ISSN 0730-0301. doi: 10.1145/3730943. URL https://doi.org/10.1145/3730943.

Rui Xiang, Rongjie Lai, and Hongkai Zhao. A dual iterative refinement method for non-rigid shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 15930–15939, 2021.

LIREN YU, Jiaming Xu, and Xiaojun Lin. Graph matching with partially-correct seeds. *J. Mach. Learn. Res.*, 22:280:1–280:54, 2021.

Silvia Zuffi and Michael J Black. Stitched puppets. In *CVPR*, 2015.

Silvia Zuffi, Angjoo Kanazawa, David W Jacobs, and Michael J Black. 3d menagerie: Modeling the 3d shape and pose of animals. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6365–6373, 2017.

APPENDICES

The main text asks whether the editor lowers error from one fixed start, where a network trained on the test family still leads, and whether a low error arrives as one piece. The sections below record the symbol table, four arguments on the three stages, wall-clock time, the knob grid, further color transfers, the contiguity counts, the published schedules of 2D-GEM and HOPE, and the DIR class diagnostic.

A NOTATION

Table 5: Notation used in the main text and appendices.

Symbol	Description	Remarks
$\mathcal{M}_s, \mathcal{M}_t$	Source / target triangle meshes	$\mathcal{M} = (\mathcal{V}, \mathbf{F})$, coordinates $\mathbf{V}$
$\mathbf{V} \in \mathbb{R}^{n\times 3}$	Vertex coordinates	$n \in \{n_s, n_t\}$
$\mathbf{p} : \mathcal{V}_t \to \mathcal{V}_s$	Hard correspondence	Target-indexed array
$\Pi \in \{0,1\}^{n_t \times n_s}$	Transport, $\Pi_{j,\mathbf{p}(j)} = 1$	Only binary map matrix used
$\mathbf{A}, \widetilde{\mathbf{A}}$	Clean / witness adjacency	Geometry mesh unchanged
$\mathbf{H}^{(k)}$	Walk-support difference	Pairs with a walk of length k and none of length $k-1$. Equals graph distance k on a clean triangle mesh; the runs build it on the witness graph
$\varepsilon(j)$	Residual at target vertex j	Equation (1); $\delta = 10^{-8}$ is only the denominator offset
$\mathbf{M}_{\mathrm{sym}} = \mathbf{D}^{-1}\mathbf{W}$	Symmetrized walk in Prop. 3	Reported runs use $\mathbf{M}_t = \mathbf{D}^{-1}\widetilde{\mathbf{A}}_t$
d_g	Intrinsic geodesic distance	Exact length in the propositions. Reported error uses d_{heat} Crane et al. (2013)
$s_k(i,j)$	k-hop overlap score	Equation (2)
$\mathcal{N}_t(j)$	Witness neighbors of j	Positive entries of row j
$\mathcal{T}$	Witness perturbation	Optional; seed 42 in the reported runs
$\mathbf{c}_s, \mathbf{c}_t$	Voronoi cluster labels	$\{1,\dots,K\}^n$; cluster index is α, not k
$\mathcal{B}$	Bad vertex set	Thresholded structural error
$\tau_{\mathrm{fix}}, q_{\mathrm{lap}}$	Structural / polish thresholds	Reported runs use 0.95
$K, k_{\max}, T$	Clusters / hop cycle / iterations	Defaults $10, 8, 32$

B THEORETICAL ANALYSIS

The statements are Propositions 1–4. This appendix records the arguments and the assumptions they actually use. $\mathcal{M}_s$ and $\mathcal{M}_t$ are connected compact piecewise-linear surfaces; G_s and G_t are their 1-skeletons. Nothing below is a recovery guarantee for the unknown ground-truth map. In this appendix d_g is exact geodesic length on that surface. Every number in the experiments uses the heat-method approximation of the same quantity.

Proposition 1 (When the residual vanishes)**.** *If* $(\Pi\widetilde{\mathbf{A}}_s\Pi^\top)[j,:] = \widetilde{\mathbf{A}}_t[j,:]$*, then* $\varepsilon(j) = 0$*.*

Proposition 2 (Cleanup freezes anchors and does not move the map)**.** *The cleanup stage does not change* $\mathbf{p}$*. For each label it freezes the largest connected component in the witness graph and gives every other vertex the label of the nearest frozen anchor set under exact geodesic length. If* a *is a closest anchor point of a reassigned vertex* v*, among all frozen anchor points, a minimizing geodesic from* v *to* a *lies in the closed cell of points at least as close to* a *as to any other anchor point. The cell of a multi-point anchor set need not be path-connected, and the new label need not be one component of the witness graph.*

Proposition 3 (One symmetrized averaging step does not raise Dirichlet energy)**.** *Let* $\mathbf{W}$ *be symmetric and nonnegative with* $D_{aa} = \sum_b W_{ab} > 0$ *for every vertex* a*, let* $\mathbf{M}_{\mathrm{sym}} = \mathbf{D}^{-1}\mathbf{W}$*, and let* $E_D(\mathbf{x}) = \frac{1}{2}\sum_{a,b} W_{ab}\|\mathbf{x}_a - \mathbf{x}_b\|_2^2$*. Then* $E_D(\mathbf{M}_{\mathrm{sym}}\mathbf{x}) \le E_D(\mathbf{x})$*.*

Proposition 4 (Greedy matching on the scores)**.** *On a bipartite graph with nonnegative edge weights, taking edges in decreasing weight and rejecting an edge that touches a used vertex is a* $\frac{1}{2}$*-approximation of the maximum-weight matching.*

B.1 WITNESS ADJACENCY

Remark 1 (Shared seed and what perturbation does not prove)**.** *Both meshes share RNG seed* 42 *before witness construction in the reported protocol. On shared-connectivity SCAPE this can yield*

identical witness graphs in index space; SCAPE ground truth is the identity, so structural methods including LOGOS benefit—remeshed comparisons are the informative ones. Perturbation breaks some exact score ties in practice; it does not *identify the true left–right sheet. Independent $\mathcal{T}$ on source and target can also reduce true-match hop overlap. Hop matrices $\mathbf{H}^{(k)}$ mark pairs with a walk of length k and none of length $k-1$; weight jitter does not change those supports. The weighted dense $\widetilde{\mathbf{A}}$* does *enter the residual (1). Chirality remains an empirical question. The heavy-tail fractions in Sec. 4.2 are not a left–right count. Table 3 also reports many-to-one collision rates: on SCAPE_r, HOPE and 2D-GEM are injective in the stored maps (*0 *collisions) while LOGOS has collision fraction* 0.201*, so the contiguity gain is not purchased by bijectivity.*

Design role. Witness perturbation is an optional protocol default. Full-split SCAPE_r ablations (Table 4) show turning it off is slightly lower than the default (0.033 ± 0.039 vs. 0.035 ± 0.045); heavy-tail fractions with no perturbation are $\mathrm{frac}{>}0.1{=}0.056$ and $\mathrm{frac}{>}0.15{=}0.042$.

B.2 Voronoi Contiguity

Proposition 5 (Open Voronoi cells of point seeds are path-connected)**.** *Let $(\mathcal{M}, d_g)$ be a connected compact piecewise-linear surface with its exact geodesic length, hence a proper geodesic length space: a minimizing geodesic exists between every pair of points. For point seeds $\sigma_1, \ldots, \sigma_K$, each open cell $C_\alpha = \{v : d_g(v, \sigma_\alpha) < d_g(v, \sigma_\beta)\ \forall \beta \neq \alpha\}$ is path-connected.*

Proof. Let $v \in C_\alpha$ and let γ be a minimizing geodesic from v to σ_α. For any point u on γ, $d_g(v, \sigma_\alpha) = d_g(v, u) + d_g(u, \sigma_\alpha)$. If some seed σ_β satisfied $d_g(u, \sigma_\beta) \leq d_g(u, \sigma_\alpha)$, the triangle inequality would give $d_g(v, \sigma_\beta) \leq d_g(v, u) + d_g(u, \sigma_\beta) \leq d_g(v, \sigma_\alpha)$, so v would not lie in the open cell C_α. Thus every point of γ lies in C_α, and the cell is path-connected. Points tied between two seeds are excluded by the strict inequality. □

The runs approximate this d_g by the heat method.

Farthest-point seeding places each new seed at a vertex of maximum distance from the seeds already chosen. It is a greedy heuristic toward a small covering radius. Cells are geometric, not semantic part labels.

B.3 Witness Refinement

Definition 1 (Witness score)**.** $s_k(i, j) = \left(\mathbf{H}_s^{(k)}[i, \mathbf{p}(\cdot)]\right)^\top \mathbf{H}_t^{(k)}[:, j]$.

Remark 2 (Heuristic status)**.** *Under seeded Erdős–Rényi models, overlap scores concentrate and separate Mao et al. (2023); YU et al. (2021). Triangle meshes are not that model. An intrinsic symmetry is another map with large hop overlap, so the score does not choose between the two sheets. We therefore use s_k as a local consensus heuristic and cycle $k = 1, \ldots, k_{\max}$ as a schedule, not as a multi-resolution theorem. If a score column is all zeros, an all-zero column that remains unmatched is set to source index* 0 *and can collide; this is rare under $\tau_{\mathrm{fix}}{=}0.95$.*

B.4 Proof of Proposition 1

The hypothesis is an equality of rows. The numerator of (1) is then zero, so $\varepsilon(j) = 0$. Entrywise, $(\Pi \widetilde{\mathbf{A}}_s \Pi^\top)_{j\ell} = \widetilde{\mathbf{A}}_s[\mathbf{p}(j), \mathbf{p}(\ell)]$, with or without injectivity. Injectivity on the star $\{j\} \cup \mathcal{N}_t(j)$ does not give the row equality: a target index outside that star can be sent to a source neighbor of $\mathbf{p}(j)$, which puts a positive weight in the transported row and a zero in the target row. A neighbor sent to a non-neighbor does the same. Two witness draws that put different directed weights on a corresponding pair do the same. The denominator only normalizes by the target row sum; $\delta = 10^{-8}$ keeps it finite and does not cancel a positive numerator. Because each directed weight is drawn independently, a perfect local isomorphism of the unweighted graphs can still have $\varepsilon(j) > 0$.

B.5 Geodesic cluster cleanup

Definition 2 (Label contiguity on a graph)**.** *A labeling is contiguous on a graph if every nonempty label induces a single connected component of that graph.*

Proof of Proposition 2. The stage reads $\mathbf{c}_t = \mathbf{c}_s \circ \mathbf{p}$ and never writes $\mathbf{p}$. Inside each label, connected components are taken in the witness adjacency, and the largest component is frozen as the anchor set A_α. Every vertex in a smaller component is assigned the label of the anchor set with smallest d_g. For fixed anchor sets, that choice minimizes the sum of distances from reassigned vertices to anchors. Let v be reassigned toward anchor set A_α, and let $a \in A_\alpha$ be a closest point of v among all frozen anchor points. Let $U = \{u : d_g(u,a) \le d_g(u,b) \text{ for every anchor point } b \ne a\}$. A minimizing geodesic from v to a stays in U. If an intermediate point u satisfied $d_g(u,b) < d_g(u,a)$ for some other anchor point b, then $d_g(v,b) \le d_g(v,u) + d_g(u,b) < d_g(v,u) + d_g(u,a) = d_g(v,a)$, so a would not be a closest anchor. Uniqueness of a is not required: a tie $d_g(u,b) = d_g(u,a)$ is still inside the closed cell. This is the point-seed argument of Proposition 5, applied to a, and it does not say that the whole set-Voronoi cell of A_α is path-connected. Two points of A_α can have disjoint neighborhoods once another anchor sits between them, so the cell of a multi-point set may have several pieces. The construction also does not claim that the new label is one component of the witness graph, nor that $\mathbf{c}_s \circ \mathbf{p}$ is contiguous on the clean 1-skeleton. That last quantity is measured after the full pipeline in Appendix H (0.872 on SCAPE_r; 0.602 on TOPKIDS). □

Among labelings that keep the surviving anchors fixed, sending each dropped vertex to the nearest anchor under d_g minimizes the sum of geodesic distances to anchors in $(\mathcal{V}_t, d_g)$. That is a property of the metric, not a claim that the anchors are the true parts.

B.6 Proof of Proposition 3

Symmetry of $\mathbf{W}$ and $D_{aa} > 0$ make $\mathbf{M}_{\text{sym}} = \mathbf{D}^{-1}\mathbf{W}$ self-adjoint for the inner product $\langle \mathbf{u}, \mathbf{v} \rangle_{\mathbf{D}} = \mathbf{u}^\top \mathbf{D}\mathbf{v}$. Write $\mathbf{M}$ for $\mathbf{M}_{\text{sym}}$ in the rest of this proof. $\langle \mathbf{u}, \mathbf{M}\mathbf{v} \rangle_{\mathbf{D}} = \mathbf{u}^\top \mathbf{W}\mathbf{v} = \langle \mathbf{M}\mathbf{u}, \mathbf{v} \rangle_{\mathbf{D}}$. Nonnegative symmetric weights put every eigenvalue λ of $\mathbf{M}$ in $[-1, 1]$, because $\frac{1}{2}\sum_{a,b} W_{ab}(x_a - x_b)^2 \ge 0$ and $\frac{1}{2}\sum_{a,b} W_{ab}(x_a + x_b)^2 \ge 0$ rearrange to $-\mathbf{x}^\top \mathbf{D}\mathbf{x} \le \mathbf{x}^\top \mathbf{W}\mathbf{x} \le \mathbf{x}^\top \mathbf{D}\mathbf{x}$. The proposition's energy sums the squared Euclidean coordinates. Each coordinate obeys the scalar quadratic form of the Laplacian $\mathbf{L} = \mathbf{D} - \mathbf{W} = \mathbf{D}(\mathbf{I} - \mathbf{M})$:

$$E_D(\mathbf{x}) = \mathbf{x}^\top \mathbf{L}\mathbf{x} = \langle \mathbf{x}, (\mathbf{I} - \mathbf{M})\mathbf{x} \rangle_{\mathbf{D}}.$$

In a $\mathbf{D}$-orthonormal eigenbasis, $\mathbf{x} = \sum_i \alpha_i \phi_i$ and $\mathbf{M}\phi_i = \lambda_i \phi_i$, so

$$E_D(\mathbf{x}) = \sum_i \alpha_i^2 (1 - \lambda_i), \qquad E_D(\mathbf{M}\mathbf{x}) = \sum_i \alpha_i^2 \lambda_i^2 (1 - \lambda_i).$$

Each coefficient satisfies $\lambda_i^2 \le 1$, and $1 - \lambda_i \ge 0$, so $E_D(\mathbf{M}\mathbf{x}) \le E_D(\mathbf{x})$. The identity $\mathbf{M}\mathbf{1} = \mathbf{1}$ is the eigenvalue $\lambda = 1$ on constants, which contributes nothing to E_D. The inequality is for this symmetric $\mathbf{W}$. The polish in the reported runs multiplies coordinates by $\mathbf{D}^{-1}\widetilde{\mathbf{A}}_t$, and the two directed weights on an edge are drawn independently, so $\widetilde{\mathbf{A}}_t$ need not be symmetric. Symmetrizing the matrix after the fact defines a different operator. Proposition 3 does not claim $E_D(\mathbf{M}_t\mathbf{x}) \le E_D(\mathbf{x})$ for that asymmetric $\mathbf{M}_t$. Replacing a smoothed point by a source vertex in (3) is not an application of $\mathbf{M}$, and it can raise E_D.

Remark 3 (Why the polish drops hop scores)**.** *The constraint $\mathbf{c}_s(i) = \mathbf{c}_t(j)$ in (3) cannot move a vertex into a different cluster. That is an invariant of the feasible set, not an optimality theorem. A long-range k-hop vote can point across a component boundary; the polish does not use s_k. Reported runs use q_{lap}=0.95 on the displacement $\rho(j) = \|\mathbf{x}(j) - \mathbf{y}(j)\|_2$ with $\mathbf{x} = \mathbf{V}_s \circ \mathbf{p}$ and $\mathbf{y} = \mathbf{M}_t\mathbf{x}$, and $\mathbf{M}_t = \mathbf{D}^{-1}\widetilde{\mathbf{A}}_t$. The update does not overwrite $\mathbf{x}$ before $\mathbf{y}$ is formed.*

B.7 Proof of Proposition 4

Let M be the matching returned by sorting nonnegative edges in decreasing weight and accepting an edge when both endpoints are free. Let $M^\star$ be any maximum-weight matching. Charge each edge $f^\star \in M^\star$ to an edge of M as follows. If $f^\star \in M$, charge it to itself. Otherwise, when $f^\star$ was considered, at least one endpoint was already covered by an accepted edge $f \in M$; charge $f^\star$ to one such f. That f was accepted earlier, so $w(f) \ge w(f^\star)$. Each vertex is incident to at most one edge of $M^\star$, so each endpoint of f receives at most one charge, and f receives at most two. Hence $w(M^\star) \le 2\, w(M)$. The witness stage runs this rule on s_k, including zero scores. An all-zero

column that remains unmatched is then set to source index 0. The Hungarian algorithm computes a maximum-weight assignment and is not this scan Kuhn (2012). That $\arg\max$ write is not covered by the $\frac{1}{2}$ statement, because it can collide. The polish (3) locks neither endpoint, so Proposition 4 does not apply to it. When $|\mathcal{B}| = \Theta(n)$ the scan is not near-linear. The threshold τ_{fix}=0.95 is how the reported runs try to keep $\mathcal{B}$ small; it is not a proof that $\mathcal{B}$ is small.

C Algorithms and Complexity

Algorithm 1 is in the main text. Initialization costs $O(n_s n_t d)$ for a dense descriptor match; the witness perturbation is $O(|E|)$; Voronoi seeding is $O(Kn \log n)$; each witness iteration builds hop shells in $O(k|E|)$ and scores $\mathcal{B}$ in $O(n|\mathcal{B}|)$, with greedy selection $O(n|\mathcal{B}| \log(n|\mathcal{B}|))$ before GPU parallelism; component cleanup is $O(n + |E|)$ plus $O(Kn \log n)$ for island geodesics; the polish is $O(T(|E| + n\bar{n}_c))$. The *current implementation* materializes dense $n \times n$ residual and score buffers ($\Theta(n^2)$ memory); the sparse bounds above are aspirational. When $|\mathcal{B}| = \Theta(n)$ the greedy step is not near-linear. The reported runs use τ_{fix}=0.95, which is intended to keep $\mathcal{B}$ small. Wall-clock times are Table 6. Those times mix two machines. LOGOS, LOGOS-lmks, LOGOS-ULRSSM-*, NAM-SHOT, NAM-ULRSSM-*, the short ZoomOut rows, and the ULRSSM rows were timed on Windows 11 Pro, Ryzen 9 9900X, 32 GB, RTX 5070 Ti 16 GB. 2D-GEM, HOPE, and plain ZoomOut were timed on an RTX 2080 Ti, i5-15300, 32 GB DDR4. The 1.35 s SCAPE_r ablation is the refiner alone; the 7.0 s LOGOS entry in Table 6 includes the rest of that pipeline.

D Full comparison

Table 1 in the main text is every stored geodesic error. Teal is the lowest number in that column and orange is the second lowest. On SCAPE_r the second place is a tie at 0.027. A dash means that row was not evaluated on that split. Plain ZoomOut is a stored spectral run whose initializer was not recorded. The sample standard deviations of the training-free LOGOS pair means are 0.005/0.044/0.023/0.013/0.013/0.059 (SCAPE, SCAPE_r, FAUST_r, TOSCA_r, TOPKIDS, SMAL_r). Table 4 reports 0.035 ± 0.045 from the leave-one-stage file. Those two standard deviations are different summaries of the SCAPE_r LOGOS row. The fractions of LOGOS pairs above 0.1 are 0/0.070/0.030/0.011/0/0.225, and above 0.15 are 0/0.056/0.003/0/0/0.077 (SCAPE, SCAPE_r, FAUST_r, TOSCA_r, TOPKIDS, SMAL_r).

Table 6: Mean seconds per pair, including the descriptor match on the SHOT row. Two machines are mixed (Appendix C), so the column is a record of wall-clock time on the machine that produced that row.

Method	SCAPE	SCAPE_r	FAUST_r	TOSCA_r	TOPKIDS	SMAL_r
SHOT k-NN	76.0	4.7	5.4	5.8	34.3	4.0
ULRSSM-faust	0.3	3.8	2.7	0.2	4.9	–
ULRSSM-scape	0.3	2.9	2.8	0.2	4.8	2.8
ULRSSM-topkids	2.1	3.6	2.9	0.4	12.6	11.5
ULRSSM-smal	–	–	–	–	–	11.2
NAM-SHOT	145.0	25.4	24.1	24.2	82.0	22.5
NAM-ULRSSM-faust	61.3	19.5	17.9	19.4	63.8	–
NAM-ULRSSM-scape	55.9	17.7	17.7	19.0	69.1	–
NAM-ULRSSM-topkids	–	27.1	26.6	–	59.9	–
NAM-ULRSSM-smal	–	–	–	–	–	24.7
ZoomOut	49.4	7.6	7.6	7.5	56.9	17.2
2D-GEM	29.0	21.2	20.4	20.3	95.0	51.4
HOPE	61.0	23.5	20.7	25.0	51.9	27.7
ZoomOut-lmks	227.3	39.3	39.5	41.4	–	48.6
ZoomOut-ULRSSM-faust	177.4	36.4	31.6	41.5	–	–
ZoomOut-ULRSSM-scape	179.7	33.2	33.4	38.6	–	35.1
ZoomOut-ULRSSM-topkids	–	–	–	–	130.2	–
ZoomOut-ULRSSM-smal	–	–	–	–	–	34.0
LOGOS-lmks	17.9	3.6	3.4	3.7	12.2	5.4
LOGOS-ULRSSM-faust	18.8	9.1	5.5	3.7	–	–
LOGOS-ULRSSM-scape	21.4	8.1	6.1	3.8	19.3	7.5
LOGOS-ULRSSM-topkids	–	–	–	–	15.7	–
LOGOS-ULRSSM-smal	–	–	–	–	–	6.8
LOGOS	55.1	7.0	6.9	6.7	30.9	10.3

Table 7: One-factor LOGOS ablations, ten pairs per dataset, seed 42, same pairs for every row of a column. Mean geodesic error (↓). The 71-pair SCAPE_r leave-one-stage table is Table 4.

Setting	SCAPE	SCAPE_r	FAUST_r	TOSCA_r	TOPKIDS	SMAL_r
Default	0.013	0.022	0.029	0.023	0.019	0.078
No perturbation	0.013	0.022	0.028	0.023	0.020	0.097
No residual gate	0.013	0.030	0.033	0.027	0.025	0.090
No component cleanup	0.013	0.026	0.029	0.026	0.021	0.074
No averaging	0.009	0.030	0.031	0.028	0.021	0.076
K=5	0.014	0.022	0.025	0.023	0.023	0.076
K=20	0.013	0.021	0.026	0.022	0.017	0.074
τ_{fix}=0.50	0.012	0.025	0.029	0.026	0.020	0.089
τ_{fix}=0.99	0.013	0.022	0.028	0.023	0.019	0.073
q_{lap}=0.50	0.060	0.024	0.028	0.024	0.024	0.078
q_{lap}=0.99	0.009	0.022	0.030	0.023	0.019	0.078
Hops 1–4	0.019	0.029	0.034	0.029	0.019	0.097
Hops 1–12	0.013	0.023	0.028	0.022	0.019	0.058

E Ablations and parameter sensitivity

Full-split leave-one-stage on SCAPE_r. Table 4 removes one stage at a time on all 71 SCAPE_r pairs and reports mean±std of the per-pair mean geodesic error. Turning the residual gate or the averaging polish off raises the mean; turning the witness perturbation off is slightly lower than the default on this remeshed set (ties broken differently, not a guarantee of a better map). The 1.35 s/pair default timing is refiner-only and is not comparable to the 7.0 s pipeline timing in Table 6. Table 4 is in the main text.

Ten-pair multi-dataset sensitivity. Additional parameter moves (K, τ_{fix}, q_{lap}, hop range) use ten pairs per dataset, drawn once without replacement from the full lists of Table 1 with seed 42 (Table 9). The same pair list is used for every setting of that dataset. The default is K=10, hops 1–8, T=32, τ_{fix}=q_{lap}=0.95, with the witness perturbation on. Table 7 reports the one-factor means. These subset numbers are for sensitivity only; Table 1 uses the full splits. SMAL_r full-split comparisons use a 1-to-1 pairing of the second correspondence column and are reported in Tables 1–6; the ten-pair study below is sensitivity only. TOSCA_r uses the same shared-reference correspondence loading as FAUST_r (per-shape files remapped through a common index), not identity.

Witness seed sensitivity. On the same ten SCAPE_r pairs, changing only the witness adjacency seed gives mean geodesic errors 0.024 (seed 0), 0.022 (seeds 1, 2, 123), and 0.022 (seed 42). The spread across seeds is small relative to the leave-one-stage gaps in Table 4. Seed 42 is the reported protocol; it is not uniquely best.

Table 8: Witness adjacency seed sensitivity on the ten SCAPE_r ablation pairs.

Seed	Mean geodesic error	Std over pairs
0	0.024	0.006
1	0.022	0.004
2	0.022	0.004
42	0.022	0.004
123	0.022	0.005

The one-factor means are Table 7.

F Additional Qualitative Results

Figures 4 and 5 are further color transfers, for SCAPE and for a ULRSSM start on FAUST_r. Figure 3 in the main text is the train/test swap. Table 2 is in the main text.

Table 9: Fixed ablation pairs (seed 42). SCAPE and SCAPE_r share the same ten pairs.

#	TOPKIDS		#	SCAPE / SCAPE_r	
	source	target		source	target
1	kid16	kid17	1	mesh000	mesh037
2	kid17	kid18	2	mesh004	mesh041
3	kid18	kid21	3	mesh010	mesh064
4	kid18	kid25	4	mesh012	mesh060
5	kid20	kid19	5	mesh018	mesh006
6	kid20	kid21	6	mesh022	mesh027
7	kid20	kid25	7	mesh033	mesh064
8	kid22	kid17	8	mesh046	mesh011
9	kid23	kid18	9	mesh049	mesh040
10	kid23	kid24	10	mesh056	mesh008

#	SMAL_r		#	FAUST_r	
	source	target		source	target
1	cougar_04	cougar_02	1	tr_reg_001	tr_reg_009
2	cougar_01	cougar_04	2	tr_reg_002	tr_reg_003
3	dog_03	dog_02	3	tr_reg_010	tr_reg_020
4	dog_02	dog_04	4	tr_reg_055	tr_reg_085
5	dog_02	dog_06	5	tr_reg_057	tr_reg_077
6	hippo_04	hippo_02	6	tr_reg_029	tr_reg_089
7	hippo_01	hippo_05	7	tr_reg_030	tr_reg_043
8	horse_09	horse_06	8	tr_reg_004	tr_reg_051
9	horse_01	horse_09	9	tr_reg_008	tr_reg_075
10	horse_02	horse_10	10	tr_reg_032	tr_reg_096

#	TOSCA_r	
	source	target
1	cat0	cat9
2	cat4	cat7
3	david0	david6
4	victoria0	victoria12
5	victoria1	victoria21
6	victoria21	victoria4
7	michael11	michael14
8	michael15	michael8
9	michael18	michael7
10	michael11	michael2

G Detecting Poor Matches

The residual $\varepsilon(j)$ in (1) decides which vertices are poorly matched. Figure 6 plots that count on one pair from FAUST_r, SCAPE_r, and TOPKIDS. The quantile rule always flags the same fraction of vertices; the fixed rule tracks how many vertices still exceed an absolute residual.

H Contiguity after refinement

Mean geodesic error can stay moderate while a cluster arrives as several islands. After LOGOS, we measure contiguity of the pushed labels $\mathbf{c}_s \circ \mathbf{p}$ on the target's clean 1-skeleton (input mesh edges, not the witness). Table 10 reports the mean number of components per nonempty cluster and the fraction of clusters that are a single component. On SCAPE_r most clusters are contiguous (0.872 single-component); on TOPKIDS topological noise leaves more islands (0.602). These numbers describe $\mathbf{c}_s \circ \mathbf{p}$ after the full pipeline; they are not a corollary of Proposition 2.

Table 11 compares SHOT nearest-neighbor initialization to LOGOS on the ten SCAPE_r ablation pairs. SHOT transfers source Voronoi labels into a highly fragmented labeling (about 145 components per nonempty cluster, none contiguous). After LOGOS the same labels average 1.17 com-

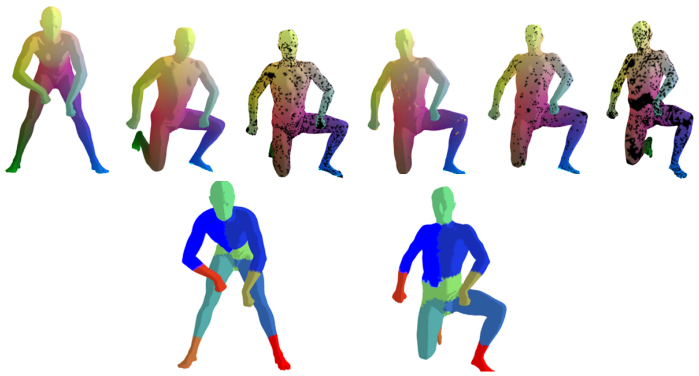

Figure 4: SCAPE qualitative comparison and LOGOS clustering (SHOT init).



Figure 5: ULRSSM+LOGOS on FAUST_r, with the Voronoi clusters used by the polish.

ponents and 0.91 of nonempty clusters are a single component. That gap is the contiguity of the pushed label field after the full refiner. The initializer alone does not produce it, and the measurement includes collisions.

Table 3 in the main text places the same metric next to HOPE, 2D-GEM, and ZoomOut on the full SCAPE_r (71) and TOPKIDS (90) splits. Counts are components of $\mathbf{c}_s \circ \mathbf{p}$ after the full pipeline.

Table 10: Contiguity of $\mathbf{c}_s \circ \mathbf{p}$ after LOGOS (clean target 1-skeleton).

Dataset	Pairs	Mean components / cluster	Fraction contiguous
SCAPE_r	71	1.25 ± 0.30	0.872
TOPKIDS	90	2.31 ± 0.96	0.602

I 2D-GEM AND HOPE

2D-GEM and HOPE are the neighborhood refiners in Table 1. Both edit a hard map. The schedules in the two papers are different from each other, and both differ from the schedule in Sec. 3.

2D-GEM initializes by a maximum-weight matching on SHOT descriptors Kamhoua et al. (2022). Each iteration runs one pass of that matching on the 2-hop adjacency, scores local map distortion on the second ring, and repairs the high-distortion vertices in a truncated Laplace–Beltrami basis. On the nearly isometric sets the paper uses 500 eigenvectors, ten iterations, and the ten distortion thresholds 0.528, 0.456, 0.384, 0.312, 0.240, and then 0.168 for the remaining five iterations. On TOPKIDS every threshold is 100, the basis is ten eigenvectors, and the 2-hop pass is repeated for 60 iterations. HOPE reads the threshold 100 as switching off that spectral repair Kamhoua & Qu (2024).

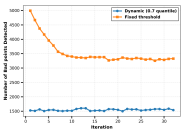

(a) FAUST_r

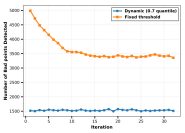

(b) SCAPE_r

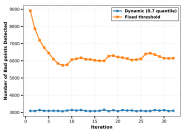

(c) TOPKIDS

Figure 6: Vertices flagged as poorly matched, on one pair from each dataset. Blue: 0.70-quantile rule (a fixed fraction of the vertices). Orange: a fixed residual threshold. The runs in Tables 1–6 use the fixed rule τ_{fix}=0.95, which flags only the tail.

Table 11: SHOT init vs LOGOS contiguity on the ten SCAPE_r ablation pairs.

Map	Mean components / cluster	Fraction contiguous
SHOT k-NN	145.4	0.00
LOGOS	1.17	0.91

HOPE keeps one schedule on every dataset. It selects poorly matched vertices by the same second-ring distortion, then matches only those vertices by k-hop overlap for $1 \le k \le k_{\max}$. The published setting is $k_{\max}$=8, 60 iterations, and ten distortion thresholds spaced from 100 down to 0.2, with the last threshold held once it is reached. The neighborhood step is therefore a longer menu of shells than 2D-GEM's 2-hop, run for more iterations than the ten used on 2D-GEM's nearly isometric sets. A variant in that paper that cycles the hops without the distortion gate is weak on the isometric shapes, so the shells and the gate are used together.

That hop difference lines up with five columns of Table 1. HOPE is lower than 2D-GEM on SCAPE (0.011 against 0.014), SCAPE_r (0.049 against 0.097), FAUST_r (0.039 against 0.087), TOSCA_r (0.030 against 0.085), and SMAL_r (0.077 against 0.187). TOPKIDS goes the other way: 2D-GEM is 0.006, LOGOS is 0.020, and HOPE is 0.029. The split on which 2D-GEM leads is the split whose published run repeats the short shell and sets the distortion thresholds to 100. Shortening LOGOS to hops 1–4 does not reproduce that 0.006: on the ten TOPKIDS pairs the mean stays 0.019, the same as the ten-pair default.

LOGOS uses the same upper shell as published HOPE, hops 1 through 8, for 32 iterations rather than 60, and only where the residual is already large. It does not change the schedule when the dataset changes, and the polish then snaps inside a cluster, which neither of those two schedules does. On every column of Table 1 the training-free LOGOS mean is lower than HOPE. Inside the LOGOS grid itself, the longer range hops 1–12 is the lowest SMAL_r ten-pair entry, 0.058 against 0.078 for hops 1–8. That longer range was not the stored protocol, and it is not the range published for 2D-GEM or for HOPE.

J DIR ON TOSCA_R, BY CLASS

DIR Xiang et al. (2021) is omitted from the main tables. Table 12 reports a per-class diagnostic of the *stored* DIR maps on TOSCA_r. The overall 0.605 is not a uniform failure. Victoria, the largest class, is 0.847; wolf, three pairs, is 0.001. The stored implementation compares geodesic rows as if they were already aligned, $|[\mathbf{G}_{S_1}]_{ij} - [\mathbf{G}_{S_2}\mathbf{P}_t]_{ij}|/\gamma$. The residual that does not assume that alignment is $|[\mathbf{G}_{S_1}]_{ij} - [\mathbf{P}_t^\top \mathbf{G}_{S_2}\mathbf{P}_t]_{ij}|/\gamma$. Fan et al. Fan et al. (2022) showed that the first form is why DIR returns bad maps. Classes whose vertex order is not a shared template are exactly where that assumption fails. We report the stored maps and do not use DIR for remeshing-collapse claims in the main text.

K LIMITATIONS

LOGOS assumes manifold triangle meshes and reliable heat-method geodesics. Anisotropic triangulations stretch combinatorial hops relative to true geodesics. We do not handle large partiality or genus changes. Greedy updates on $\mathcal{B}$ do not enforce a global bijection. Left–right flips are not

Table 12: DIR on TOSCA_r, mean geodesic error by class (stored maps; diagnostic only).

Class	Pairs	Geodesic error	Seconds / pair
wolf	3	0.001	4.6
dog	36	0.361	8.6
horse	28	0.441	10.5
cat	55	0.495	10.8
centaur	15	0.586	11.6
michael	60	0.661	11.4
david	21	0.710	12.0
victoria	66	0.847	12.0
Overall	284	0.605	11.0

solved. The heavy-tail fractions in Sec. 4.2 are not a left–right count. The shared-seed witness protocol can make witness graphs identical in index space on shared-connectivity SCAPE. Runtime comparisons mix Python/CUDA and MATLAB stacks on different GPUs; Table 6 is not like-for-like. The current implementation materializes dense $n \times n$ buffers. LOGOS-lmks is worse than SHOT-initialized LOGOS on the shared-connectivity check (0.044 against 0.010) and is the lowest complete row on SMAL_r (0.033 against 0.056).